

Transferring Obligations Through Synchronizations

Jafar Hamin 

imec-DistriNet, Department of Computer Science, KU Leuven, Belgium
jafar.hamini@cs.kuleuven.be

Bart Jacobs 

imec-DistriNet, Department of Computer Science, KU Leuven, Belgium
bart.jacobs@cs.kuleuven.be

Abstract

One common approach for verifying safety properties of multithreaded programs is assigning appropriate permissions, such as ownership of a heap location, and obligations, such as an obligation to send a message on a channel, to each thread and making sure that each thread only performs the actions for which it has permissions and it also fulfills all of its obligations before it terminates. Although permissions can be transferred through synchronizations from a sender thread, where for example a message is sent or a condition variable is notified, to a receiver thread, where that message or that notification is received, in existing approaches obligations can only be transferred when a thread is forked. In this paper we introduce two mechanisms, one for channels and the other for condition variables, that allow obligations, along with permissions, to be transferred from the sender to the receiver, while ensuring that there is no state where the transferred obligations are lost, i.e. where they are discharged from the sender thread but not loaded onto the receiver thread yet. We show how these mechanisms can be used to modularly verify deadlock-freedom of a number of interesting programs, such as some variations of *client-server* programs, *fair readers-writers locks*, and *dining philosophers*, which cannot be modularly verified without such transfer. We also encoded the proposed separation logic-based proof rules in the VeriFast program verifier and succeeded in verifying the mentioned programs.

2012 ACM Subject Classification Theory of computation → Separation logic; Software and its engineering → Deadlocks; Software and its engineering → Process synchronization; Software and its engineering → Formal software verification; Theory of computation → Hoare logic

Keywords and phrases Hoare logic, separation logic, modular program verification, synchronization, transferring obligations, deadlock-freedom.

Digital Object Identifier 10.4230/LIPIcs.ECOOP.2019.22

Funding This research received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731453 (project VESSEDA), as well as Flemish Research Fund project grant G.0962.17N.

Acknowledgements We thank three anonymous reviewers for their careful reading of our manuscript and their many insightful comments and suggestions, and also Amin Timany for his guidance on Coq.

1 Introduction

One common approach for verifying safety properties of multithreaded programs, such as absence of data races and deadlock, is assigning appropriate permissions [3, 37] and obligations [29, 33, 1] to each thread and making sure that each thread only performs the actions for which it has permissions and it also fulfills all of its obligations before it terminates. In a separation logic-based approach [42], for example, the ownership of a heap location is a permission for accessing that location, which is assigned to the thread allocating that location. Since there is only one instance of such permission, only one thread, the one having



© Jafar Hamin and Bart Jacobs;

licensed under Creative Commons License CC-BY

33rd European Conference on Object-Oriented Programming (ECOOP 2019).

Editor: Alastair F. Donaldson; Article No. 22; pp. 22:1–22:58

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

22:2 Transferring Obligations Through Synchronizations

such permission, can access that location, which ensures absence of data races. Absence of deadlock, as another example, can be verified by making sure that for any waitable object o , such as a lock, a channel, or a condition variable, for which a thread is waiting there is a thread obliged to fulfill an obligation for that object (which only waits for objects whose levels, some arbitrary numbers¹ associated with waitable objects, are lower than the level of o) [33, 2, 24, 15]. In this setting a thread can discharge an obligation for a lock, a channel, or a condition variable by releasing that lock, sending a message on that channel, or notifying that condition variable (under some conditions), respectively.

Permissions can be transferred through synchronizations [33, 34] by consuming them in the sender thread, where for example a message is sent or a condition variable is notified, and producing them in the receiver thread, where that message or that condition variable is received. However, this technique cannot be used for transferring obligations because the transferred obligations are lost if no thread receives them. Additionally, even in the presence of a receiver there might be a state where these obligations are discharged from the sender but not received by the receiver yet. If the obligations are transferred through channels, for example, if the `receive` is not scheduled to be executed *immediately* after the `send` then there would be a state where no thread holds the transferred obligations, which makes the approach unsound. Note that in the case that the `send` operation is synchronous [29], i.e. the sender thread is suspended until a thread receives the sent message, sending obligations through channels is perfectly fine because these obligation are never lost.

In this paper we introduce two mechanisms allowing threads to transfer their obligations through synchronizations while ensuring that there is no state where these obligations are lost, i.e. where they are discharged from the sender thread but not loaded onto the receiver thread yet. The main idea behind the first mechanism is that when the transferred obligations are discharged from the sender thread the levels of these obligations have been already loaded onto a (receiver) thread. These levels are discharged from the receiver thread when it receives the transferred obligations. In the second mechanism, which is specifically for condition variables, the obligations are discharged from a notifying thread only if there is a waiting thread which is notified and receives these obligations. We show that using these mechanisms in some modular approaches, verifying finite blocking [2] and deadlock-freedom of channels and locks [33, 24], semaphores [21], and monitors [15], enables them to verify a wider range of interesting programs, such as some variations of *client-server* programs, *fair mutexes*, *fair readers-writers locks*, and *dining philosophers*. We encoded the proposed proof rules in the VeriFast program verifier [25, 26, 22] and succeeded in verifying the programs above². Additionally, we proved the soundness of both mechanisms (see Appendixes C and D): the soundness proof of the second mechanism is machine-checked with Coq³.

In the rest of this paper Sections 2 and 3 introduce two mechanisms for transferring obligations through channels and notifications, Section 4 discusses related work, and Section 5 draws a conclusion.

¹ In this paper, for simplicity we use numbers as levels, but any partially ordered set can be used as the set of levels.

² The proof of these programs, verified by the VeriFast program verifier, can be found in [17].

³ The soundness proof of the second mechanism, machine-checked in Coq, can be found in [17].

2 Transferring Obligations Through Channels

In this section we first review an approach, introduced by Leino *et al.* [33], which modularly verifies deadlock-freedom of channels. Then we extend this approach such that it also allows obligations to be transferred from the thread sending on a channel to the thread receiving from that channel. Lastly, we provide some variations of a client-server program which can be verified thanks to such extension.

2.1 Verifying Channels

Leino *et al.* [33] introduced a modular approach to verify deadlock-freedom of programs which communicate through channels. The main idea in this approach is to associate with each thread a bag⁴ of channels, namely the bag of *obligations* of that thread, which must be empty when that thread terminates. A thread can discharge an obligation for a channel by either sending a message on that channel or delegating that obligation to a forked thread. This approach guarantees absence of deadlock by making sure that for any thread trying to receive a message from a channel ch there is either a message in ch or an obligation for ch in the bag of obligations of a thread which only waits for objects whose levels are lower than the level of ch (preventing circular dependencies). This constraint is applied by making sure that if a thread tries to receive from a channel ch then 1) it spends a *credit* for ch , where a thread can obtain a credit for ch by adding ch to the bag of its obligations, and 2) the level of ch is lower than the levels of the obligations of that thread. Note that a credit for a channel in this approach indicates that there is either a message on that channel or there is a thread holding an obligation for that channel.

As an example, consider the deadlock-free program shown in the left hand side of Figure 1, where after creating a channel ch , the main thread first forks a receiver thread, trying to receive a message from ch , and then sends a message on ch . The verification of this program is shown in the right hand side of this figure, where we use separation logic [42] to reason about the ownership of permissions. The verification of this program is started with an empty bag of obligations, denoted by $\text{obs}(\{\})$. As indicated below each command, by creating a channel a (*duplicable* and *leakable*) permission channel for accessing that channel is produced and an arbitrary level, denoted by a function R , is assigned to that channel. Before forking the receiver thread, using a *ghost command*⁵ g_credit , a credit and an obligation for this channel are loaded onto the main thread. The former is given to the forked thread, where this credit is spent by the command $\text{receive}(ch)$, and the latter is discharged by the main thread when it executes the command $\text{send}(ch, 12)$.

The separation logic-based proof rules, introduced by Jacobs *et al.* [23, 24], used to verify this program are shown in Figure 2. Note that the postcondition of each command in this figure is a lambda expression that given the return value of that command returns the assertion held after execution of that command. When a channel is created, as shown in Rule NEWCHANNEL, any arbitrary level can be assigned to that channel by the proof author. Note that, generally, this level must be chosen such that the constraint number 2, mentioned above, is met at each receive operation. Sending a message on a channel, as shown in Rule SEND, discharges an obligation for that channel. As shown in Rule RECEIVE, a thread can

⁴ We model bags of objects as functions from objects to natural numbers. We also use \uplus indicating the union (i.e. the pointwise sum) of two bags.

⁵ The ghost commands are inserted into the program for verification purposes and have no effect on the program's behavior.

```

ch := new_channel();
fork(
  receive(ch);
  send(ch, 12)
);

```

```

{obs({})}
ch := new_channel();
{obs({}) * channel(ch) ∧ R(ch)=r}
g_credit(ch);
{obs({ch}) * channel(ch) * credit(ch)}
fork(
  {obs({}) * channel(ch) * credit(ch) ∧ ch < {} }
  receive(ch)
  {obs({})}
);
{obs({ch}) * channel(ch)}
send(ch, 12)
{obs({})}

```

■ **Figure 1** Verification of deadlock-freedom of channels, where O in $\text{obs}(O)$ is the bag of obligations of the running thread, $R(ch)$ denotes the level of ch , and $ch < O \Leftrightarrow \forall o \in O. R(ch) < R(o)$.

try to receive a message from a channel ch only if that thread spends a credit for ch and the level of ch is lower than the levels of all obligations of that thread. As shown in Rule FORK, a thread can transfer a part of its permissions and obligations to a forked thread, provided that the forked thread discharges all the delegated obligations. Lastly, as shown in Rule CREDIT, using a ghost command g_credit a thread can obtain a credit for a channel if that channel is loaded onto the bag of the obligations of that thread.

It can be proved that any program verified by the mentioned proof rules, where the verification starts from an empty bag of obligations and also ends with such bag, never deadlocks, i.e. it always has a running thread, not waiting for any channel, until it terminates. We know that for any channel ch all of these proof rules preserve the following invariant, where $Ct(ch)$ and $Ot(ch)$ denote the total number of credits and obligations for ch in the system, respectively, and $\text{size}(ch)$ denotes the number of messages in ch :

$$Ct(ch) \leq Ot(ch) + \text{size}(ch) \quad (1)$$

Now consider a deadlocked state, where each thread of a verified program is waiting for a channel. Among all of these channels take the one having a minimal wait level, namely ch_{min} . Since $\text{size}(ch_{min})=0$ and there exists a credit for ch_{min} in the system held by the waiting thread, according to the invariant above and the constraint number 2, there exists a thread having an obligation for ch_{min} that is waiting for a channel whose level is lower than the level of ch_{min} , which contradicts minimality of the level of ch_{min} .

2.2 Transferring Permissions and Obligations Through Channels

The approach presented in the previous section fails to verify some applications of channels where some obligations must be transferred from a thread sending on a channel to the thread receiving from that channel. Consider the *client-server* program shown in Figure 3, for example, where the server waits to receive a message that consists of a client request, which must be processed by the server, and a client channel, from which the client expects to receive the response of the server. Although the client routine in this example is deadlock-free, the Leino *et al.* approach fails to verify this program, since the routine `main` cannot give any credit for the channel ch' , which is created inside the client, to the client. This program can be verified if after creating ch' a credit and an obligation for ch' are loaded onto the client

$$\begin{array}{c}
\text{NEWCHANNEL} \\
\{\text{true}\} \text{ new_channel } \{\lambda ch. \text{ channel}(ch) \wedge R(ch)=r\} \\
\\
\text{SEND} \\
\{\text{obs}(O) * \text{ channel}(ch)\} \text{ send}(ch, m) \{\lambda_. \text{ obs}(O - \{ch\}) * \text{ channel}(ch)\} \\
\\
\text{RECEIVE} \\
\{\text{obs}(O) * \text{ channel}(ch) * \text{ credit}(ch) \wedge ch \prec O\} \text{ receive}(ch) \{\lambda_. \text{ obs}(O) * \text{ channel}(ch)\} \\
\\
\text{FORK} \\
\frac{\{a * \text{ obs}(O)\} c \{\lambda_. \text{ obs}(\{\})\}}{\{a * \text{ obs}(O \uplus O')\} \text{ fork}(c) \{\lambda_. \text{ obs}(O')\}} \\
\\
\text{CREDIT} \\
\{\text{obs}(O) * \text{ channel}(ch)\} \text{ g_credit}(ch) \{\lambda_. \text{ obs}(O \uplus \{ch\}) * \text{ channel}(ch) * \text{ credit}(ch)\}
\end{array}$$

■ **Figure 2** Proof rules ensuring deadlock-freedom of channels

<pre> routine server(channel ch){ (req, ch') := receive(ch); result := process(req); send(ch', result) </pre>	<pre> routine client(channel ch){ ch' := new_channel(); send(ch, (request(), ch')); receive(ch') </pre>	<pre> routine main(){ ch := new_channel(); fork(server(ch)); client(ch) </pre>
--	--	---

■ **Figure 3** A client-server program

such that the latter is transferred to the server through the client's request and the former is spent for the command `receive(ch')` executed in the client.

Similar to permissions, two necessary conditions for transferring obligations are: 1) when a thread sends a message on a channel ch the transferred obligations of ch , which are transferred through ch , are discharged from the bag of the obligations of that thread, and 2) when a thread receives a message from a channel ch the transferred obligations of ch are loaded onto the bag of the obligations of that thread. However, these conditions are not sufficient because these obligations are lost if no thread receives from ch . Additionally, even in the presence of a receiver if the `receive(ch)` is not scheduled to be executed *immediately* after the `send(ch, m)` then there would be a state where no thread holds the transferred obligations, which makes the approach unsound.

To address this problem, in addition to the bag of obligations, we associate with each thread a new bag, namely the bag of *importers* of that thread, which consists of the channels that transfer (import) some obligations to that thread. Similar to the bag of obligations, the bag of importers of a thread must be empty when that thread terminates and a thread can wait for a channel ch only if the level of ch is lower than the levels of all obligations which are *possibly* imported by all importers of that thread except for ch itself. Having this bag, we enforce an additional condition, numbered 3, when a thread sends some obligations on an importer channel, which imports some obligations, this channel must be already in the bag of importers of a (receiver) thread. This importer channel is discharged from the receiver thread as it receives a message from this channel. This additional condition is met by making sure that any thread sending on an importer channel ch spends a *transferring credit* for ch , where a thread can obtain a transferring credit of ch by adding ch to the bag

22:6 Transferring Obligations Through Synchronizations

of its importers.

Formally, the third condition, which holds for any importer channel ch , ensures that for any transferring credit of ch or any message in ch (which means a transferring credit of ch has been spent by sending a message on ch and no thread has received it yet), there exists an instance of ch in the bag of importers of a (receiver) thread. This invariant is shown in the following as Invariant 2, where $M^r(ch)$ denotes the bag of the levels of the obligations which are *possibly* imported by ch , $Tt(ch)$ denotes the total number of transferring credits of ch in the system, $It(ch)$ denote the number of occurrences of ch in the bags of importers of all threads in the system, and $size(ch)$ denotes the number of messages in the queue of ch .

$$\forall ch. M^r(ch) \neq \emptyset \Rightarrow Tt(ch) + size(ch) \leq It(ch) \quad (2)$$

Additionally, the two first conditions mentioned above, as well as the ones mentioned in the previous section, ensure that if a thread waits on a channel ch' then there is either an obligation of ch' in the system, or a message on ch' , or a message in the queue of a channel through which an obligation of ch' is transferred (which means an obligation of ch' has been transferred through this message and no thread has received it yet). This invariant is formally shown in the following as Invariant 3, where $Ot(ch)$ and $Ct(ch)$ denote the total number of obligations and credits of ch in the system, $queue(ch)$ denotes the list of messages in the channel ch , and $M'(ch)$ is a function, that given a message, specifies the bag of the obligations which are transferred through that message in ch . Note that the levels of these obligations must be in the bag of the levels of the obligations which are possibly imported by ch , as shown formally in Invariant 4, where $levels(O)$ maps each element of O to its level.

$$\forall ch'. Ct(ch') \leq Ot(ch') + size(ch') + \sum_{ch} \sum_{m \in queue(ch)} M'(ch)(m)(ch') \quad (3)$$

$$\forall ch, m \in queue(ch). levels(M'(ch)(m)) \subseteq M^r(ch) \quad (4)$$

It can be proved that any program preserving such invariants never deadlocks, i.e. it always has a running thread, not waiting for any channel, until it terminates. Consider a deadlocked state, where each thread of a program is waiting for a channel. Among all of these channels take the one having a minimal wait level, namely ch_{min} . Since $size(ch_{min})=0$ and there exists a credit for ch_{min} in the system held by the waiting thread, according to Invariant 3, there exists either 1) a thread having an obligation of ch_{min} that is waiting for a channel whose level is lower than the level of ch_{min} , or 2) there exists a message m on a channel ch through which an obligation of ch_{min} is transferred, i.e. $0 < M'(ch)(m)(ch_{min})$ which by Invariant 4 implies $R(ch_{min}) \in M^r(ch)$, where $R(ch_{min})$ denotes the level of ch_{min} . In the first case minimality of the level of ch_{min} is contradicted. In the second case, since $0 < size(ch)$ and $M^r(ch) \neq \emptyset$, by Invariant 2, there exists a thread having an importer channel ch that is waiting for a channel whose level is lower than the level of ch_{min} (because ch imports an obligation of level of ch_{min}), which again contradicts minimality of ch_{min} .

The proof rules enforcing such invariants are shown in Figure 4, where $M^r(ch)$ denotes the bag of the levels of the obligations which are possibly imported by ch , and the parameters M and M' in the permission **channel** of a channel are two functions that given a message return the permissions and the obligations which are transferred through that message. When a channel ch is created, as shown in Rule **NEWCHANNEL**, the value of these functions for this channel can be specified arbitrarily. As shown in Rule **SEND**, when a message m is sent on ch , the permissions which are transferred through m as well as one transferring credit of

$$\begin{array}{l}
\text{NEWCHANNEL} \\
\{\text{true}\} \text{ new_channel } \{\lambda ch. \text{ channel}(ch, M, M') \wedge R(ch)=r \wedge M^r(ch)=R\} \\
\\
\text{SEND} \\
\{\text{obs}(O, I) * \text{ channel}(ch, M, M') * M(m) * (M^r(ch)=\{\!\!\!\!\!\} \vee \text{trandit}(ch)) \wedge \\
\text{ levels}(M^r(m)) \subseteq M^r(ch)\} \text{ send}(ch, m) \\
\{\lambda _. \text{ obs}(O - \{ch\} - M^r(m), I) * \text{ channel}(ch, M, M')\} \\
\\
\text{RECEIVE} \\
\{\text{obs}(O, I) * \text{ channel}(ch, M, M') * \text{ credit}(ch) \wedge ch \prec O \wedge ch \prec^r I\} \text{ receive}(ch) \\
\{\lambda m. \text{ obs}(O \uplus M^r(m), I - \{ch\}) * \text{ channel}(ch, M, M') * M(m)\} \\
\\
\text{FORK} \\
\frac{\{a * \text{ obs}(O, I)\} c \{\lambda _. \text{ obs}(\{\!\!\!\!\!\}, \{\!\!\!\!\!\})\}}{\{a * \text{ obs}(O \uplus O', I \uplus I')\} \text{ fork}(c) \{\lambda _. \text{ obs}(O', I')\}} \\
\\
\text{CREDIT} \\
\{\text{obs}(O) * \text{ channel}(ch, M, M')\} \text{ g_credit}(ch) \\
\{\lambda _. \text{ obs}(O \uplus \{ch\}) * \text{ channel}(ch, M, M') * \text{ credit}(ch)\} \\
\\
\text{TRANDIT} \\
\{\text{obs}(O, I) * \text{ channel}(ch, M, M')\} \text{ g_trandit}(ch) \\
\{\lambda _. \text{ obs}(O, I \uplus \{ch\}) * \text{ channel}(ch, M, M') * \text{ trandit}(ch)\}
\end{array}$$

■ **Figure 4** The updated proof rules ensuring deadlock-freedom of importer channels, where $M^r(ch)$ denotes the bag of the levels of the obligations which are possibly imported by ch ; M and M' in the permission `channel` of a channel are functions, that given a message, return the permissions and the obligations which are transferred through that message, respectively; bag I in the permission `obs`(O, I) of a thread denotes the channels importing some obligations onto that thread; `levels`(O) maps each element of O to its level; and $ch \prec^r I \Leftrightarrow \forall ch' \in I. ch = ch' \vee ch \prec M^r(ch')$.

ch , denoted by `trandit`(ch), (if ch is an importer channel) are consumed and the obligations which are transferred through m as well as an obligation of ch are discharged from the bag of the obligations. Additionally, this rule makes sure by adding m to the queue of ch Invariant 4 is still preserved. As shown in Rule `RECEIVE`, when a thread tries to receive a message m from ch the level of ch must be lower than the levels of the obligations of the thread and also the levels of the obligations which are possibly imported by all importers of that thread except for ch itself, i.e. $\forall ch' \in I. ch = ch' \vee ch \prec M^r(ch')$ where I is the bag of the importers of the receiving thread. Additionally, a credit of ch is consumed, the permissions which are transferred through m are produced, ch is discharged from the bag of the importers, and the obligations which are transferred through m are loaded onto the bag of the obligations. As shown in Rule `FORK`, a thread can transfer a part of its permissions, obligations, and importer channels to a forked thread, provided that the forked thread discharges all of the delegated obligations and importer channels. As shown in Rule `TRANDIT` a thread can obtain a transferring credit of ch by loading ch onto the bag of its importers.

The verification of the program in Figure 3 using the proposed proof rules is illustrated in Figure 5. Note that for the sake of readability we elide repeated occurrences of the permissions `channel` in the proof of the given programs. As shown in the precondition and the postcondition of the routine `server`, denoted by `req` and `ens`, this routine discharges an

importer channel ch , where a message from ch is received. Since this routine tries to wait on ch , it requires a credit and a permission channel for ch . The permission channel in the precondition of this routine indicates that along with the client's request the server receives an obligation and a permission channel for the client's channel through which no permission or obligation is transferred. The specification of the routine `client` indicates that this routine discharges an obligation for the server channel ch , since it sends a request to this channel. As it is shown in the verification of this routine, after creating the client channel ch' and before sending it to the server, one obligation and one credit for ch' are loaded onto the system. The former is transferred to the server, where it is discharged by sending a response on ch' , and the latter is spent in the rest of this routine for receiving from ch' . The routine `main` in this program is successfully verified, since starting with an empty bag of obligations and importers the verification of this routine is finished with such bags too.

2.3 Conditional Channels

A different variation of a client-server program is shown in Figure 6, where a server keeps accepting the clients' requests until it receives a specific message `done`. The verification of this program using the proposed proof rules is illustrated in Figure 7. Note that if the client sends a message `done` to the server the session is finished and the client should not transfer an obligation of ch' through this message, i.e. an obligation of ch' is transferred through a message sent on ch only if that message is not a message `done`. Since in this program a client can send multiple requests to the server and for each request it requires a permission `trandit(ch)`, the server sends this permission to the client each time it replies to the client, i.e. a permission `trandit(ch)` is transferred through the client's channel ch' . Additionally, since the server might wait for ch more than once and for each wait it requires a `credit(ch)`, before replying to the client a credit and an obligation for ch are loaded onto the server and the loaded obligation is transferred to the client, where this obligation is discharged by sending on ch , i.e. an obligation for ch is transferred through the client's channel ch' . Since the channel ch' transfers some obligations from the server to the client, the server requires a permission `trandit(ch')` before sending on ch' , which can be obtained from the client through the channel ch , i.e. if the client sends a request (and not a message `done`) to the server it also sends a permission `trandit(ch')` to the server.

2.4 Server Channels

Another variation of a client-server program is shown in Figure 8, where a server *infinitely* accepts the clients' requests through a *server channel* s . A desired property of these kinds of programs is that if they have a thread waiting for a *non-server* channel they also have a running thread, not waiting for any channel. In other words, a program can be considered as a terminated program even if there are still some specific threads, namely *daemon threads*, such as a server thread or a garbage collector thread, which are not terminated. In Java these threads are terminated by the JVM when there is no longer any user thread running.

To achieve the mentioned property we only need to make sure that if a thread tries to receive from a server channel s it has no obligation and the bag of the importers of this thread only contains s .

It can be proved that any program verified by enforcing the constraint above meets the mentioned desired property, i.e. if this program has a thread waiting for a non-server channel it also has a running thread, not waiting for any channel. Consider an undesired state, where each thread is waiting for a channel while some of these channels are non-server channels.


```

Mch ::=  $\lambda m. \text{channel}(\text{snd}(m), \lambda \_ . \text{true}, \lambda \_ . \{\}) \wedge M^r(\text{snd}(m)) = \{\}$ 
M'ch ::=  $\lambda m. \{\text{snd}(m)\}$ 

routine server(channel ch){
req : {obs( $\{\}, \{ch\}$ ) * channel(ch, Mch, M'ch) * credit(ch)}
  (req, ch') := receive(ch);
  {obs( $\{ch'\}, \{\}$ ) * channel(ch',  $\lambda \_ . \text{true}, \lambda \_ . \{\}) \wedge M^r(ch') = \{\}$ }
  result := process(req);
  send(ch', result)
ens : {obs( $\{\}, \{\}$ )}

routine client(channel ch){
req : {obs( $\{ch\}, \{\}$ ) * channel(ch, Mch, M'ch) * trandit(ch)  $\wedge M^r(ch) = \{1\}$ }
  ch' := new_channel();
  {obs( $\{ch\}, \{\}$ ) * channel(ch',  $\lambda \_ . \text{true}, \lambda \_ . \{\})$  * trandit(ch)  $\wedge R(ch') = 1 \wedge M^r(ch') = \{\}$ }
  g_credit(ch');
  {obs( $\{ch, ch'\}, \{\}$ ) * trandit(ch) * credit(ch')}
  send(ch, (request(), ch'));
  {obs( $\{\}, \{\}$ ) * credit(ch')}
  receive(ch')
ens : {obs( $\{\}, \{\}$ )}

routine main(){
req : {obs( $\{\}, \{\}$ )}
  ch := new_channel();
  {obs( $\{\}, \{\}$ ) * channel(ch, Mch, M'ch)  $\wedge R(ch) = 1 \wedge M^r(ch) = \{1\}$ }
  g_credit(ch);
  {obs( $\{ch\}, \{\}$ ) * credit(ch)}
  g_trandit(ch);
  {obs( $\{ch\}, \{ch\}$ ) * credit(ch) * trandit(ch)}
  fork(
    {obs( $\{\}, \{ch\}$ ) * credit(ch)}
    server(ch)
    {obs( $\{\}, \{\}$ )};
  {obs( $\{ch\}, \{\}$ ) * trandit(ch)}
  client(ch)
ens : {obs( $\{\}, \{\}$ )}

```

■ **Figure 5** Verification of the program in Figure 3

<pre> routine cserver(channel ch){ (req, ch') := receive(ch); while(req \neq done){ send(ch', process(req)); (req, ch') := receive(ch) } </pre>	<pre> routine client(channel ch){ ch' := new_channel(); send(ch, (request(), ch')); receive(ch'); send(ch, (done, ch')) </pre>	<pre> routine main() { ch := new_channel(); fork(cserver(ch)); client(ch) </pre>
---	---	---

■ **Figure 6** A server keeps serving a client until receiving a specific message done

22:10 Transferring Obligations Through Synchronizations

```

Mch(ch) ::= λm. channel(snd(m), λ_. trandit(ch), λ_. {ch}) ∧ Mr(snd(m))={1} *
fst(m)=done ? true : trandit(snd(m))
M'ch ::= λm. fst(m)=done ? {} : {snd(m)}

routine cserver(channel ch){
req : {obs({}, {ch}) * channel(ch, Mch(ch), M'ch) * credit(ch) ∧ R(ch)=1 ∧ Mr(ch)={1}}
(req, ch') := receive(ch);
{obs(req=done ? {} : {ch'}, {}) * channel(ch', λ_. trandit(ch), λ_. {ch}) *
(req=done ? true : trandit(ch')) ∧ Mr(ch')={1}}
while(req ≠ done){
inv : {Mr(ch')={1} ∧ req=done ? obs({}, {}) : (obs({ch'}, {}) * trandit(ch'))}
  g_trandit(ch); g_credit(ch);
  {obs({ch', ch}, {ch}) * trandit(ch') * trandit(ch) * credit(ch)}
  send(ch', process(req));
  {obs({}, {ch}) * credit(ch)}
  (req, ch') := receive(ch)
} ens : {obs({}, {})}

routine client(channel ch){
req : {obs({ch}, {}) * channel(ch, Mch(ch), M'ch) * trandit(ch) ∧ R(ch)=1 ∧ Mr(ch)={1}}
ch' := new_channel();
{obs({ch}, {}) * trandit(ch) * channel(ch', λ_. trandit(ch), λ_. {ch}) ∧ R(ch')=1 ∧
Mr(ch')={1}}
g_credit(ch'); g_trandit(ch');
{obs({ch, ch'}, {ch'}) * trandit(ch) * credit(ch') * trandit(ch')}
send(ch, (request(), ch'));
{obs({}, {ch'}) * credit(ch')}
receive(ch');
{obs({ch}, {}) * trandit(ch)}
send(ch, (done, ch'))
ens : {obs({}, {})}

routine main(){
req : {obs({}, {})}
ch := new_channel();
{obs({}, {}) * channel(ch, Mch(ch), M'ch) ∧ R(ch)=1 ∧ Mr(ch)={1}}
g_credit(ch) ; g_trandit(ch);
{obs({ch}, {ch}) * credit(ch) * trandit(ch)}
fork(
  {obs({}, {ch}) * credit(ch)}
  cserver(ch)
  {obs({}, {})});
{obs({ch}, {}) * trandit(ch)}
client(ch)
ens : {obs({}, {})}

```

■ **Figure 7** Verification of the program in Figure 6, where $a ? b : c$ evaluates to b if the value of a is true, and otherwise to c .

```

routine server(channel s){
  while(true){
    (req, ch') := receive(s);
    send(ch', process(req))
  }
}

routine client(channel s){
  ch' := new_channel();
  send(s, (request(), ch'));
  receive(ch')
}

routine main(){
  s := new_channel();
  fork(server(s));
  fork(client(s));
  client(s)
}

```

■ **Figure 8** A server keeps serving clients

```

NEWCHANNEL
{true} new_channel {λch. channel(ch, M, M') ∧ R(ch)=r ∧ Mr(ch)=R ∧ S(ch)=b}

RECEIVE
{obs(O, I) * channel(ch, M, M') * (S(ch) ∨ credit(ch)) ∧ ch < O ∧ ch <r I ∧
  (¬S(ch) ∨ (O=∅ ∧ ∀ch'∈I. ch'=ch))} receive(ch)
{λm. obs(O⊕M'(m), I-∅{ch}) * channel(ch, M, M') * M(m)}

TRANDITS
  {obs(O, I) * channel(ch, M, M')} g_trandits(ch)
  {λ_. obs(O, I⊕{ch∞}) * channel(ch, M, M') * trandit∞(ch)}

```

■ **Figure 9** The updated proof rules ensuring deadlock-freedom of server channels, where S is a function that given a channel specifies whether that channel is a server channel or not, and o^∞ represents an infinite number of occurrences of o .

Among all of these non-server channels take the one having a minimal wait level, namely ch_{min} . Since $\text{size}(ch_{min})=0$ and there exists a credit for ch_{min} in the system held by the waiting thread, by Invariant 3, either 1) there exists a thread having an obligation of ch_{min} that is waiting for a channel ch whose level is lower than the level of ch_{min} , or 2) there exists a message m in a channel ch through which an obligation of ch_{min} is transferred, i.e. $0 < M'(ch)(m)(ch_{min})$ which by Invariant 4 implies $R(ch_{min}) \in M^r(ch)$, where $R(ch_{min})$ denotes the level of ch_{min} . The first case contradicts minimality of the level of ch_{min} , because in this case ch is a non-server channel, since that thread is waiting for ch while it has some obligations. In the second case, since $0 < \text{size}(ch)$ and $M^r(ch) \neq \emptyset$, by Invariant 2, there exists a thread t having an importer channel ch that is waiting for a channel ch_1 whose level is lower than the level of ch_{min} (because ch imports an obligation of level of ch_{min}). Since $0 < \text{size}(ch)$ we know $ch_1 \neq ch$. Accordingly, ch_1 cannot be a server channel (because t is trying to receive from ch_1 while it has an importer channel ch which is not equal to ch_1), which is a contradiction.

The proof rules which need to be updated are shown in Figure 9. As shown in Rule NEWCHANNEL, when a channel is created it must be specified whether this channel is a server channel or not, which is denoted by a function S . As shown in Rule RECEIVE, if a thread tries to receive from a server channel ch the bag of the obligations of this thread must be empty and the bag of the importer channels of this thread must only contain ch . Note that a thread does not need to spend a credit for waiting on a server channel. We also introduce a new ghost command $g_trandits(ch)$, shown in Rule TRANDITS, which produces an infinite number of transferring credits of ch , denote by $trandit^\infty(ch)$, by loading an infinite number of ch , denoted by ch^∞ , onto the importers. The verification of the program in Figure 8 using these rules is shown in Figure 10, where the server does not need any credit

22:12 Transferring Obligations Through Synchronizations

```

Ms(s) ::= λm. channel(snd(m), λ_. trandit(s), λ_. {} ) ∧ Mr(snd(m))={ }
M's ::= λm. {snd(m)}

routine server(channel s){
req : {obs({ }, {s∞}) * channel(s, Ms(s), M's) ∧ Mr(s)={1} ∧ S(s)}
while(true){
inv : {obs({ }, {s∞})}
  (req, ch') := receive(s);
  {obs({ch'}, {s∞}) - {s}} * channel(ch', λ_. trandit(s), λ_. {} ) ∧ Mr(ch')={ }
  g_trandit(s);
  {obs({ch'}, {s∞}) * trandit(s)}
  send(ch', process(req))
  {obs({ }, {s∞})}
}
ens : {false}}

routine client(channel s){
req : {obs({ }, { }) * channel(s, Ms(s), M's) * trandit(s) ∧ Mr(s)={1}}
ch' := new_channel();
{obs({ }, { }) * trandit(s) * channel(ch', λ_. trandit(s), λ_. {} ) ∧ R(ch')=1 ∧ Mr(ch')={ }}
g_credit(ch');
{obs({ch'}, { }) * trandit(s) * credit(ch')}
send(s, (request(), ch'));
{obs({ }, { }) * credit(ch')}
receive(ch')
{obs({ }, { }) * trandit(s)}
ens : {obs({ }, { })}}

routine main(){
req : {obs({ }, { })}
s := new_channel();
{obs({ }, { }) * channel(s, Ms(s), M's) ∧ Mr(s)={1}}
g_trandits(s);
{obs({ }, {s∞}) * trandit∞(s)}
fork(
  {obs({ }, {s∞})}
  server(s);
  {obs({ }, { }) * trandit∞(s)}
fork(client(s);
  {obs({ }, { }) * trandit∞(s)}
  client(s)
ens : {obs({ }, { })}}

```

■ **Figure 10** Verification of the program in Figure 8

```

routine server(channel s){
  while(true){
    (thr, ch') := receive(s);
    ch := new_channel();
    send(ch', (through, ch));
    fork(cserver(ch))
  }}

routine client(channel s){
  ch' := new_channel();
  send(s, (through, ch'));
  (thr, ch) := receive(ch');
  send(ch, (request(), ch'));
  (res, ch) := receive(ch');
  send(ch, (done, ch'))}

routine main()
{
  s := new_channel();
  fork(server(s));
  fork(client(s));
  client(s)
}

```

■ **Figure 11** A server keeps serving clients through separate conditional channels

for waiting on the server channel s . Note that this verification ensures neither termination nor deadlock-freedom. It actually ensures that if this program has a thread waiting for the non-server channel ch' it also has a running thread, not waiting for any channel.

Using the proposed proof rules it is also possible to verify some other variations of client-server programs such as the one shown in Figure 11, where a server infinitely serves clients' request through separate conditional channels (see Appendix A illustrating the proof of this program). Note that the definition of the function `cserver` in this figure is similar to the one shown in Figure 6.

3 Transferring Obligations Through Notifications

In this section we introduce a mechanism which allows obligations to be transferred from a thread notifying a condition variable (CV) to the notified thread. The main idea behind this mechanism is based on this unique feature of condition variables that when a sender thread notifies a CV, if there is a receiver thread waiting for that CV, the receiver thread *immediately* receives this notification. Accordingly, if there exists a thread waiting for a condition variable v , it is safe to discharge the transferred obligations of v when v is notified and load these obligations to the receiver thread as it is notified. Note that a notification on a condition variable is lost if there is no thread waiting for that condition variable.

The number of threads waiting for a condition variable can be tracked by using the approach introduced by Hamin *et al.* [15, 16], which modularly verifies deadlock-freedom of programs synchronized by monitors. Since this approach only allows permissions to be transferred from a notifying thread to the one notified, it cannot verify some interesting programs such as a particular implementation of fair readers-writers locks, shown in Figure 19, and dining philosophers (see Appendix B). In the following we first review this approach and then we extend this approach such that it also allows transferring of obligations, enabling it to verify a wider range of programs such as the ones we just mentioned.

3.1 Verifying Monitors

Hamin *et al.* [15, 16] introduced a modular approach for verifying deadlock-freedom of programs synchronized by condition variables (CVs), where executing a command `wait(v, l)` on a CV v , which is associated with a lock l , releases l and suspends the running thread, and executing a command `notify(v)/notifyAll(v)` wakes up one/all thread(s) waiting for CV v , if any. This approach ensures absence of deadlock by making sure that for any CV v for which a thread is waiting there is a thread obliged to fulfill an obligation for v which only waits for waitable objects whose levels are lower than the level of v . In this approach when a thread acquires a lock l , the total number of waiting threads, and the total number of obligations of any CV v associated with l , denoted by $Wt(v)$ and $Ot(v)$ respectively, can be mentioned in

22:14 Transferring Obligations Through Synchronizations

$$\begin{array}{c}
\text{NEWLOCK} \\
\{\text{true}\} \text{ newlock } \{\lambda l. \text{ ulock}(l, \{\}, \{\}) \wedge R(l)=r\} \\
\\
\text{INITLOCK} \\
\{\text{ulock}(l, Wt, Ot) * \text{inv}(Wt, Ot) * \text{obs}(O)\} \text{ g_initl}(l) \{\lambda _. \text{lock}(l) * \text{obs}(O) \wedge I(l)=\text{inv}\} \\
\\
\text{ACQUIRE} \\
\{\text{lock}(l) * \text{obs}(O) \wedge l \prec O\} \text{ acquire}(l) \\
\{\lambda _. \exists Wt, Ot. \text{locked}(l, Wt, Ot) * I(l)(Wt, Ot) * \text{obs}(O \uplus \{l\})\} \\
\\
\text{RELEASE} \\
\{\text{locked}(l, Wt, Ot) * I(l)(Wt, Ot) * \text{obs}(O \uplus \{l\})\} \text{ release}(l) \{\lambda _. \text{lock}(l) * \text{obs}(O)\} \\
\\
\text{NEWCV} \\
\{\text{true}\} \text{ new_cvar } \{\lambda v. \text{ucond}(v) \wedge R(v)=r\} \\
\\
\text{INITCV} \\
\{\text{ucond}(v) * \text{ulock}(l, Wt, Ot)\} \text{ g_initc}(v) \{\lambda _. \text{cond}(v, M) * \text{ulock}(l, Wt, Ot) \wedge L(v)=l\} \\
\\
\text{WAIT} \\
\{\text{cond}(v, M) * \text{locked}(l, Wt, Ot) * I(l)(Wt \uplus \{v\}, Ot) * \text{obs}(O \uplus \{l\}) \wedge \\
l=L(v) \wedge v \prec O \wedge l \prec O \wedge \text{enoughObs}(v, Wt \uplus \{v\}, Ot)\} \text{ wait}(v, l) \\
\{\lambda _. \text{cond}(v, M) * \text{obs}(O \uplus \{l\}) * \exists Wt', Ot'. \text{locked}(l, Wt', Ot') * I(l)(Wt', Ot') * M\} \\
\\
\text{NOTIFY} \\
\{\text{cond}(v, M) * \text{locked}(L(v), Wt, Ot) * (Wt(v)=0 \vee M)\} \text{ notify}(v) \\
\{\lambda _. \text{cond}(v, M) * \text{locked}(L(v), Wt - \{v\}, Ot)\} \\
\\
\text{NOTIFYALL} \\
\{\text{cond}(v, M) * \text{locked}(L(v), Wt, Ot) * \left(\begin{array}{c} Wt(v) \\ * \\ M \end{array} \right)_{i=1}\} \text{ notifyAll}(v) \\
\{\lambda _. \text{cond}(v, M) * \text{locked}(L(v), Wt[v:=0], Ot)\} \\
\\
\text{CHARGE OBLIGATION} \\
\{\text{obs}(O) * \text{ulock/locked}(L(v), Wt, Ot)\} \text{ g_chrg}(v) \\
\{\lambda _. \text{obs}(O \uplus \{v\}) * \text{ulock/locked}(L(v), Wt, Ot \uplus \{v\})\} \\
\\
\text{DISCHARGE OBLIGATION} \\
\{\text{obs}(O) * \text{ulock/locked}(L(v), Wt, Ot) \wedge \text{enoughObs}(v, Wt, Ot - \{v\})\} \text{ g_disch}(v) \\
\{\lambda _. \text{obs}(O - \{v\}) * \text{ulock/locked}(L(v), Wt, Ot - \{v\})\}
\end{array}$$

■ **Figure 12** Proof rules verifying deadlock-freedom of monitors, where $Wt(v)$ and $Ot(v)$ denote the total number of threads waiting for v and the total number of obligations for v , respectively; the parameter M in the permission cond of a condition variable denotes the permissions which are transferred from the thread notifying that condition variable to the one(s) notified; $L(v)$ denotes the lock associated with the condition variable v ; $\text{enoughObs}(v, Wt, Ot) \Leftrightarrow (Wt(v) > 0 \Rightarrow Ot(v) > 0)$; and $v \prec O \Leftrightarrow \forall o \in O. R(v) < R(o)$.

the proof of that thread. In order to ensure the mentioned constraint this approach makes sure that 1) if a command $\text{wait}(v, l)$ is executed then $0 < Ot(v)$, i.e. there is an obligation of v in the system, 2) if an obligation of v is discharged then after this discharge the invariant $0 < Wt(v) \Rightarrow 0 < Ot(v)$ holds, i.e. if there is a thread waiting for v then after this discharge there are still some obligations for v in the system, and 3) a thread executes a command $\text{wait}(v, l)$ only if the level of v is lower than the levels of the obligations of that thread.

A program in this approach can be successfully verified if each lock associated with some CVs has an appropriate invariant such that for any CV v associated with that lock this invariant implies $0 < Wt(v) \Rightarrow 0 < Ot(v)$. Accordingly, in this approach each lock invariant is parametrized over the bags Wt and Ot , which map all CVs associated with that lock to the number of their waiting threads and obligations, respectively.

The proof rules proposed in this approach are shown in Figure 12. As shown in Rule NEWLOCK, when a lock l is created an arbitrary level is assigned to that lock and an uninitialized lock permission $\text{ulock}(l, \{\}, \{\})$ is produced. The second and the third parameters of this permission are two bags mapping the CVs associated with l to their number of waiting threads and obligations, respectively. As shown in Rule INITLOCK, this uninitialized lock permission can be converted to a (*duplicable and leakable*) lock permission $\text{lock}(l)$ if the assertion resulting from applying the invariant of that lock, denoted by $l(l)$, to the bags stored in the permission ulock is consumed (the permissions described by the invariant of this lock are transferred from the thread to the lock). As shown in Rule ACQUIRE, when a thread acquires this lock the permissions described by the invariant of this lock are transferred from the lock to the thread. Additionally, a permission $\text{locked}(l, Wt, Ot)$ is provided for the thread, where Wt and Ot are two bags mapping the CVs associated with l to their number of waiting threads and obligations, respectively, and are existentially quantified in the postcondition. Note that to prevent circular dependencies the precondition of this rule enforces that the level of l be lower than the levels of the obligations of the acquiring thread. Additionally, this lock is added to the bag of the obligations of this thread. As shown in rule RELEASE, when this lock is released it is discharged from the bag of the obligations and the assertion resulting from applying the invariant of this lock to the bags stored in the permission locked is consumed. Additionally, the permission locked is converted to a permission lock .

As shown in Rule NEWCV, when a CV is created an arbitrary level is assigned to it and an uninitialized permission ucond for that CV is produced. As shown in Rule INITCV, this permission can be converted to a (*duplicable and leakable*) permission cond if a lock is associated to this CV, denoted by $L(v)$. Additionally, the transferred permissions of this CV, which are transferred from the notifying thread to the one notified, are also specified in this rule, denoted by M in the permission cond . These permissions are consumed when a command $\text{notify}(v)$ is executed (if there is a thread waiting for v ; see the precondition of Rule NOTIFY), and are produced when a command $\text{wait}(v, l)$ is executed (see the postcondition of Rule WAIT). Note that $\text{notifyAll}(v)$ transfers $Wt(v)$ instances of these permissions, denoted by $\prod_{i=1}^{Wt(v)} M$ (see the precondition of Rule NOTIFYALL). As shown in Rule WAIT, when a command $\text{wait}(v, l)$ is executed, since l is going to be released and the number of threads waiting for v is going to be increased, the result of applying the invariant of lock l to bags $Wt \uplus \{v\}$ and Ot must be consumed, where Wt and Ot are the bags stored in the permission locked of l . Additionally, the level of v must be lower than the levels of all obligations of the thread except for l . Note that the level of l must be lower than the levels of these obligations, too, since when the thread is woken up it tries to reacquire l . As previously mentioned, the precondition of this rule also makes sure that $0 < Ot(v)$, which is enforced by the invariant $\text{enoughObs}(v, Wt \uplus \{v\}, Ot)$. This invariant follows from $l(l)(Wt \uplus \{v\}, Ot)$

22:16 Transferring Obligations Through Synchronizations

$$\begin{array}{l}
\text{INITCV} \\
\{ \text{ucond}(v) * \text{unlock}(l, Wt, Ot) \} \text{g_initc}(v) \\
\{ \lambda _ . \text{cond}(v, M, M') * \text{unlock}(l, Wt, Ot) \wedge \text{L}(v)=l \} \\
\\
\text{WAIT} \\
\{ \text{cond}(v, M, M') * \text{locked}(l, Wt, Ot) * \text{l}(l)(Wt \uplus \{v\}, Ot) * \text{obs}(O \uplus \{l\}) \wedge \\
\quad \text{l}=\text{L}(v) \wedge v \prec O \wedge \text{l} \prec O \uplus M' \wedge \text{enoughObs}(v, Wt \uplus \{v\}, Ot) \} \text{wait}(v, l) \\
\{ \lambda _ . \text{cond}(v, M, M') * \text{obs}(O \uplus \{l\} \uplus M') * \exists Wt', Ot'. \text{locked}(l, Wt', Ot') * \text{l}(l)(Wt', Ot') * M \} \\
\\
\text{NOTIFY} \\
\{ \text{obs}(O \uplus (0 < Wt(v) ? M' : \emptyset)) * \text{cond}(v, M, M') * \text{locked}(\text{L}(v), Wt, Ot) * (Wt(v)=0 \vee M) \} \\
\text{notify}(v) \{ \lambda _ . \text{obs}(O) * \text{cond}(v, M, M') * \text{locked}(\text{L}(v), Wt - \{v\}, Ot) \} \\
\\
\text{NOTIFYALL} \\
\{ \text{cond}(v, M, \emptyset) * \text{locked}(\text{L}(v), Wt, Ot) * \left(\begin{array}{c} Wt(v) \\ * \\ i:1 \end{array} M \right) \} \text{notifyAll}(v) \\
\{ \lambda _ . \text{cond}(v, M, \emptyset) * \text{locked}(\text{L}(v), Wt[v:=0], Ot) \}
\end{array}$$

■ **Figure 13** New proof rules verifying deadlock-freedom of monitors allowing transferring obligations through notifications, where the parameter M' in the permission `cond` of a condition variable denotes the obligations which are transferred from the thread notifying that condition variable to the one notified.

if the invariant of l is properly defined such that for any CV v' associated with l and any Wt' and Ot' , we have $\text{l}(l)(Wt, Ot) \Rightarrow \text{enoughObs}(v', Wt, Ot)$. Lastly the precondition of this command makes sure that v is associated with lock l , which is enforced by $\text{L}(v)=l$. As shown in Rules NOTIFY/NOTIFYALL, when a CV v is notified, one/all instance(s) of v is/are removed from the bag Wt stored in the permission `locked` of the lock associated with v , if any. Unlike the Leino *et al.* approach [33], this approach has no notion of credits and an obligation for a CV v is loaded/unloaded when that obligation is also loaded/unloaded onto/from the bag Ot stored in the permission `locked` of the lock associated with v , as shown in Rules CHARGE OBLIGATION and DISCHARGE OBLIGATION. However, an obligation for v is discharged only if after this discharge we have $0 < Wt(v) \Rightarrow 0 < Ot(v)$, which is enforced by the invariant `enoughObs` in the precondition of the rule DISCHARGE OBLIGATION.

3.2 Transferring Obligations Through Notifications

In this subsection we extend the Hamín *et al.* approach [15] such that it also allows obligations to be transferred from the notifying thread to the one notified. To this end we make sure that 1) when a thread notifies a CV v the transferred obligations of v , which are transferred through a notification on v , are discharged from the thread only if there is a thread waiting for v which is notified, that is $0 < Wt(v)$, 2) when a thread waits for a CV v the transferred obligations of v are loaded onto the bag of obligations of that thread as that thread is notified, and 3) when a thread waits for a CV v the level of the lock associated with v is lower than the levels of the transferred obligations of v , too, since when the thread wakes up, where these obligations are loaded onto the bag of the obligations of the thread, this lock must be reacquired. For the sake of simplicity, any CV v on which `notifyAll(v)` is performed must have no transferred obligations. The proof rules which need to be updated are shown in Figure 13.


```

routine new_mutex(){
  l := new_lock;
  v := new_cvar;
  mutex(l:=l, v:=v,
    b:=new_bool(false))
}

routine enter_cs(mutex m){
  acquire(m.l);
  while(m.b)
    wait(m.v, m.l);
  m.b := true;
  release(m.l)}

routine exit_cs(mutex m){
  acquire(m.l);
  m.b := false;
  notify(m.v);
  release(m.l)}

```

■ **Figure 14** Mutexes

```

routine main(){
  m := new_mutex();
  fork(
    while(1){
      enter_cs(m); /* CS */
      exit_cs(m)});
  enter_cs(m); /* CS */
  exit_cs(m)}

```

■ **Figure 15** The main thread is starved if it is scheduled only when the forked thread is in the CS.

It can be proved that any program verified by the mentioned proof rules, where the verification starts from an empty bag of obligations and also ends with such bag, never deadlocks, i.e. it always has a running thread, not waiting for any waitable object such as a condition variable or a lock, until it terminates. We know that for any waitable object o all of these proof rules preserve the invariant $0 < Wt(o) \Rightarrow 0 < Ot(o)$, where $Wt(o)$ and $Ot(o)$ denote the total number of waiting threads and obligations for o in the system, respectively. Note that this invariant holds even when some obligations are transferred because these obligations are immediately transferred from the notifying thread to the one notified. Now consider a deadlocked state, where each thread of a verified program is waiting for an object. Among all of these objects take the one having a minimal wait level, namely o_{min} . By the invariant above, there exists a thread having an obligation for o_{min} that is waiting for an object whose level is lower than the level of o_{min} , which contradicts minimality of the level of o_{min} .

```

routine new_mutex(){
  l := new_lock;
  v := new_cvar;
  mutex(l:=l, v:=v,
    b:=new_bool(false),
    w:=new_int(0))
}

routine enter_cs(mutex m){
  acquire(m.l);
  if(m.b){
    m.w := m.w+1;
    wait(m.v, m.l)}
  else
    m.b := true;
  release(m.l)}

routine exit_cs(mutex m){
  acquire(m.l);
  m.b := false;
  if(0 < m.w){
    m.w := m.w-1;
    m.b := true;
    notify(m.v)}
  release(m.l)}

```

■ **Figure 16** Fair mutexes on top of fair monitors

22:18 Transferring Obligations Through Synchronizations

$$\text{mutex}(\text{mutex } m, \text{waitobj } o) = \text{lock}(m.l) * \text{cond}(m.v, \text{true}, \{m.v\}) \wedge \\ l(m.l)=\text{linv}(m) \wedge L(m.v)=m.l \wedge R(m.l) < R(m.v) \wedge o=m.v$$

$$\text{linv}(\text{mutex } m) = \\ \lambda Wt. \lambda Ot. \exists b, w. m.b \mapsto b * m.w \mapsto w \wedge Wt(m.v)=w \wedge (b ? 0 < Ot(v) : Wt(v) = 0)$$

```

routine new_mutex(){
req : {true}
  l := new_lock;
  {unlock(l, {}, {}) ∧ R(l)=r-1}
  v := new_cvar; g_initc(v);
  {unlock(l, {}, {}) * cond(v, true, {v}) ∧ R(v)=r ∧ L(v)=l}
  m := mutex(l:=l, v:=v, b:=new_bool(false), w:=new_int(0)); g_initl(l); m
ens : {λm. ∃o. mutex(m, o) ∧ R(o)=r}

```

```

routine enter_cs(mutex m){
req : {obs(O) * mutex(m, o) ∧ o < O}
  acquire(m.l);
  {obs(O ⊕ {m.l}) * mutex(m, o) * ∃ Wt, Ot. locked(m.l, Wt, Ot) * linv(m)(Wt, Ot)}
  if(m.b){
    m.w := m.w+1;
    {obs(O ⊕ {m.l}) * mutex(m, o) * ∃ Wt, Ot. locked(m.l, Wt, Ot) * linv(m)(Wt ⊕ {m.v}, Ot)}
    wait(m.v, m.l)
    {obs(O ⊕ {m.l, m.v}) * mutex(m, o) * ∃ Wt, Ot. locked(m.l, Wt, Ot) * linv(m)(Wt, Ot)}
  }
  else{
    m.b := true; g_chrg(m.v)
    {obs(O ⊕ {m.l, m.v}) * mutex(m, o) * ∃ Wt, Ot. locked(m.l, Wt, Ot) * linv(m)(Wt, Ot)}
  };
  {obs(O ⊕ {m.l, m.v}) * mutex(m, o) * ∃ Wt, Ot. locked(m.l, Wt, Ot) * linv(m)(Wt, Ot)}
  release(m.l)
ens : {obs(O ⊕ {o}) * mutex(m, o)}

```

```

routine exit_cs(mutex m){
req : {obs(O ⊕ {o}) * mutex(m, o) ∧ o < O}
  acquire(m.l);
  {obs(O ⊕ {m.v, m.l}) * mutex(m, o) * ∃ Wt, Ot. locked(m.l, Wt, Ot) * linv(m)(Wt, Ot)}
  m.b := false;
  if(0 < m.w){
    m.w := m.w-1; m.b := true; notify(m.v)
    {obs(O ⊕ {m.l}) * mutex(m, o) * ∃ Wt, Ot. locked(m.l, Wt, Ot) * linv(m)(Wt, Ot)}
  }
  else{
    g_disch(m.v)
    {obs(O ⊕ {m.l}) * mutex(m, o) * ∃ Wt, Ot. locked(m.l, Wt, Ot) * linv(m)(Wt, Ot)};
  }
  release(m.l)
ens : {obs(O) * mutex(m, o)}

```

■ **Figure 17** Verification of the fair mutexes implementation shown in Figure 16

3.3 Fair Mutexes

In this section we show how our extension helps to verify a fair implementation of a mutex, in which threads are synchronized by CVs. Before introducing this implementation, consider a simple (unfair) implementation of a mutex, shown in Figure 14. In this implementation a mutex consists of a boolean variable b , indicating whether the critical section (CS) is executed by any thread or not, a lock l , protecting this variable from concurrent accesses, and a CV v , preventing threads from entering the CS if there is a thread running that CS. As shown in the routine `enter_cs`, when a thread tries to enter a CS, protected by a mutex m , it first acquires the mutex's lock and while there is a thread running the CS it releases that lock and suspends itself. If that thread is notified (and reacquires the mutex's lock) while there is no thread running the CS it changes the value of the variable b to `true`, preventing other threads from entering the CS, and releases the mutex's lock. Before leaving the critical section, as shown in routine `exit_cs`, this thread acquires the mutex's lock, changes the value of the variable b to `false`, allowing other threads to enter the CS, notifies the condition variable of the mutex, waking up a waiting thread, if any, and finally releases the mutex's lock.

However, one problem with this implementation is that a thread might be in *starvation*; it might infinitely wait for entering the CS. For example, consider the program in Figure 15, where the main thread is starved if it is scheduled only when the forked thread is in the CS. In this situation, when the forked thread exits the CS, since $m.b = \text{false}$, this thread without waiting for $m.v$ again changes $m.b$ to `true` and enters the CS. To address this problem a new variable w can be added to the structure of the mutex, tracking the number of threads waiting for that mutex. Introducing this variable, the operations `enter_cs` and `exit_cs` can be updated, as shown in Figure 16. As shown in the routine `enter_cs(m)`, when a thread tries to enter a CS, if the CS is currently executed by another thread ($m.b = \text{true}$) this thread increases the variable $m.w$ and waits for a notification on $m.v$. Otherwise, this thread changes $m.b$ to `true` and continues its execution. As shown in the routine `exit_cs(m)`, when a thread leaves a CS it first changes $m.b$ to `false` and if there is a thread waiting for this CS it decreases the number of the waiting threads, changes $m.b$ to `true` and notifies a thread waiting for $m.v$. Having this implementation, the forked thread in Figure 15 cannot enter the critical section if the main thread is already waiting to enter. Note that this mutex is fair under the assumption that the monitor primitives are fair, i.e. the lock and condition variable wait queues are FIFO.

This implementation can be verified against the expected specifications, shown in Figure 17, if an obligation of $m.v$ is transferred from the thread leaving the CS to the next thread entering the CS. This transfer is necessary because one of the desired invariants of such program is $m.b \Rightarrow 0 < Ot(m.v)$, that is if the CS is executed by any thread there exists an obligation of $m.v$ in the system. Since in the fair implementation of `exit_cs` before notifying $m.v$ the variable $m.b$ is changed to `true`, an obligation of $m.v$ must be loaded onto the system. This obligation can be transferred to the notified thread, which is going to enter the CS. Note that this transfer is sound since it is only performed when $0 < m.w$, that is there is a thread waiting for $m.v$ which immediately receives the transferred obligation.

3.4 Fair Readers-Writers Locks

In this section we show how our extension makes it possible to verify a fair implementation of a readers-writers lock, which is synchronized by CVs. Before that, consider a naive implementation of a readers-writers lock (writers-preference), shown in Figure 18, which can be verified by the Hamin *et al.* [15] approach. This lock consists of three variables

22:20 Transferring Obligations Through Synchronizations

```

routine new_rdwr(){
  l := new_lock();
  vw := new_cvar;
  vr := new_cvar;
  rdwr(l:=l, vw:=vw, vr:=vr,
    aw := new_int(0), ww := new_int(0),
    ar := new_int(0))}

routine acquire_write(rdwr b){
  acquire(b.l);
  while(b.aw+b.ar>0){
    b.ww := b.ww+1;
    wait(b.vw, b.l);
    b.ww := b.ww-1 };
  b.aw := b.aw+1;
  release(b.l)}

routine release_write(rdwr b){
  acquire(b.l);
  b.aw := b.aw-1;
  notify(b.vw);
  if(b.ww = 0)
    notifyAll(b.vr);
  release(b.l)}

routine acquire_read(rdwr b){
  acquire(b.l);
  while(b.aw+b.ww>0)
    wait(b.vr, b.l);
  b.ar := b.ar+1;
  release(b.l)}

routine release_read(rdwr b){
  acquire(b.l);
  b.ar := b.ar-1;
  notify(b.vw);
  release(b.l)}

routine main(){
  rw := new_rdwr();
  fork(
    while(1){
      acquire_read(rw); /* reading ... */
      release_read(rw)
    });
  while(1){
    acquire_write(rw); /* writing ... */
    release_write(rw)
  }}

```

■ **Figure 18** Readers-writers locks synchronized by condition variables

aw , ww , and ar , keeping track of the total number of active writers, waiting writers, and active readers respectively, a lock l , protecting these variables from concurrent accesses, a condition variable v_w , preventing writers from writing while other threads are reading or writing, and a condition variable v_r , preventing readers from reading while there is another thread writing or waiting to write.

However, in this implementation due to the same reason mentioned in the previous section a writer might be starved by other writers, where for example a writer continuously releases the writing lock and without waiting for the CV v_w immediately reacquires it.

A solution to solve this problem is that when a writer releases the writing lock, if there is a waiting writer, in addition to notifying that waiting writer, it also increases the number of active writers. In other words, if there is a waiting writer in the system, the number of the active writers is increased in advance by the thread releasing the writing lock (or the reading lock) and not by the thread acquiring the writing lock. A fair implementation of the readers-writers lock (writers-preference) following this idea along with the details of the specifications of each routine is shown in Figure 19⁶. When a readers-writers lock b is created, a permission $rw(b, O_w, O_r)$ is provided, where O_w and O_r are two bags of waitable objects whose levels are in an arbitrary client-specified range R . A thread can acquire a writing lock of b if the levels of objects in O_w are lower than the levels of all obligations of that thread. When that lock is acquired the objects in O_w are loaded onto the bag of the obligations of

⁶ We inserted `assert(e)` commands, shorthand for `while($\neg e$){}` (loop forever if e is false), to simplify the proof. They can be eliminated using ghost state [27]. The verification of this program without using the `assert` commands can be found in [17].

```

rw(rdwr b, bag⟨waitobj⟩ Ow, bag⟨waitobj⟩ Or) =
lock(b.l) * cond(b.vr, true, {})* cond(b.vw, true, {b.vw}) ∧ l(ch.l)=linv(b) ∧
R(b.l) < R(b.vw) < R(b.vr) ∧ Ow={b.vr, b.vw} ∧ Or={b.vw}

linv(rdwr b) =
λ Wt. λ Ot. ∃ ar, aw, ww. b.ar ↦ ar * b.aw ↦ aw * b.ww ↦ ww ∧ L(b.vr)=L(b.vw)=b.l ∧
0 ≤ ar ∧ 0 ≤ aw ∧ 0 ≤ ww ∧ Wt(b.vw)=ww ∧
ar + aw ≤ Ot(b.vw) ∧ (Wt(b.vw) = 0 ∨ 0 < ar + aw) ∧
aw + ww ≤ Ot(b.vr) ∧ (Wt(b.vr) = 0 ∨ 0 < aw + ww)

routine new_rdwr()
req : {R ≅< Q}
ens : {λb. ∃Ow, Or. rw(b, Ow, Or) ∧
levels(Ow) ⊆ R ∧ levels(Or) ⊆ R}
{l := new_lock();
vw := new_cvar; g_initc(vw);
vr := new_cvar; g_initc(vr);
b := rdwr(l:=l, vw:=vw, vr:=vr,
aw:=new_int(0), ww:=new_int(0),
ar:=new_int(0)); g_initl(l);
b}

routine acquire_write(rdwr b)
req : {obs(O) * rw(b, Ow, Or) ∧ Ow < O}
ens : {obs(O ⊕ Ow) * rw(b, Ow, Or)}
{acquire(b.l); g_chrg(b.vr);
if(b.aw+b.ar>0){
b.ww := b.ww+1;
wait(b.vw, b.l)
}
else{ b.aw := b.aw+1; g_chrg(b.vw);
release(b.l)}}

routine release_write(rdwr b)
req : {obs(O ⊕ Ow) * rw(b, Ow, Or) ∧ Ow < O}
ens : {obs(O) * rw(b, Ow, Or)}
{acquire(b.l);
assert(0 < b.aw);
b.aw := b.aw-1;
if(b.ww > 0){
notify(b.vw);
b.ww := b.ww-1;
b.aw := b.aw+1
}
else{ notifyAll(b.vr); g_disch(b.vw);
g_disch(b.vr);
release(b.l)}}

routine acquire_read(rdwr b)
req : {obs(O) * rw(b, Ow, Or) ∧ Or < O}
ens : {obs(O ⊕ Or) * rw(b, Ow, Or)}
{acquire(b.l);
while(b.aw+b.ww>0)
wait(b.vr, b.l);
b.ar := b.ar+1; g_chrg(b.vw);
release(b.l)}}

routine release_read(rdwr b)
req : {obs(O ⊕ Or) * rw(b, Ow, Or) ∧ Or < O}
ens : {obs(O) * rw(b, Ow, Or)}
{acquire(b.l);
assert(0 < b.ar);
b.ar := b.ar-1;
if(b.ar = 0 ∧ b.ww > 0){
notify(b.vw);
b.ww := b.ww-1;
b.aw := b.aw+1
} else g_disch(b.vw);
release(b.l)}}

routine main()
req : {obs({})}
ens : {obs({})}
{rw := new_rdwr();
fork(
while(1){
acquire_read(rw); /* reading ... */
release_read(rw)
}
);
while(1){
acquire_write(rw); /* writing ... */
release_write(rw)
}
}

```

■ **Figure 19** Fair readers-writers locks on top of fair monitors, where $R \cong_{<} \mathbb{Q}$ indicates that R is order-isomorphic to the rational numbers.

this thread, which are discharged when this thread releases this lock. Similar to a writing lock, a reading lock of b can be acquired if the levels of the objects in O_r are lower than the levels of the obligations of the reading thread. When this lock is acquired the objects in O_r are loaded which are discharged when this lock is released. Note that in this program one of the desired invariants is $b.ar + b.aw \leq Ot(b.v_w)$. Accordingly, since the variable tracking the number of active writers (aw) is increased in the thread notifying v_w and not in the notified one, it is necessary to load an obligation of v_w onto the notifying thread and then transfer this obligation to the notified thread, through the notification.

4 Related Work

Permissions

One common approach to prove safety properties of a program is assigning permissions to the threads of that program and to make sure that each thread only performs actions for which that thread has permissions [3]. This approach has been adopted by concurrent separation logic [37], where a thread can access a heap location only if it owns that location. This logic has been extended to handle dynamic thread creation [11, 19, 13], rely/guarantee [48, 8], reentrant locking [12], and channels [18, 40].

Jung *et al.* [27] proposed a concurrent separation logic, namely *Iris*, for reasoning about safety of concurrent programs, as the logic in logical relations, and to reason about type systems and data abstraction, among other things. In this logic user-defined protocols on shared state are expressed through *partial commutative monoids* and are enforced through *invariants*. However, the *Iris* program logic and many other logics such as [41, 36] only prove per-thread safety (i.e. no thread ever crashes): their adequacy theorems state that the program does not reach a state where some thread cannot make a step. This works because these logics do not consider blocking constructs, where a thread may legitimately be stuck temporarily. It follows that these logics do not support programs that use primitive blocking operations. Recently, an extension of *Iris*, namely *Iron* [1], exploits a notion of obligation to prove absence of resource leaks, but not deadlock-freedom. The adequacy theorem of this logic only considers the state reached by a program after it is completely finished (i.e. all threads have reduced to a value), and it proves that in that state all resources have been freed.

Deadlock

Several approaches to verify termination [35, 14, 43], total correctness [4], and lock-freedom [20] of concurrent programs have been proposed. These approaches are only applicable to non-blocking algorithms, where the suspension of one thread cannot lead to the suspension of other threads. Consequently, they cannot be used to verify deadlock-freedom of programs using condition variables or channels, where the suspension of a notifying/sending thread might cause a waiting thread to be infinitely blocked. In [39] a compositional approach to verify termination of multi-threaded programs is introduced, where *rely-guarantee reasoning* is used to reason about each thread individually while there are some assertions about other threads. In this approach a program is considered to be terminating if it does not have any infinite computations. As a consequence, it is not applicable to programs using condition variables because a waiting thread that is never notified cannot be considered as a terminating thread. There are some other works on verifying deadlock-freedom and starvation-freedom of concurrent objects with partial methods, which do not return under certain circumstances

such as acquiring a held lock [34]. In addition to locks these approaches allows to verify other concurrent objects such as sets, stacks and queues. However, this approach is not applicable to condition variables because of *lost notifications*, i.e. a notification on a condition variable is lost if no thread is waiting for that condition variable. Note that releasing a lock, pushing an item into stack, and enqueueing an item when there is no thread waiting for the related concurrent object is not lost, since the next thread, which tries to acquire/pop/dequeue the concurrent object, will not be blocked.

There are also some other approaches addressing some common synchronization bugs of programs in the presence of condition variables. In [49], for example, an approach to identify some potential problems of concurrent programs using condition variables is presented. However, it does not take the order of execution of these commands into account. In other words, it might accept an undesired execution trace where the waiting thread is scheduled after the notifying thread, that might lead the waiting thread to be infinitely suspended. [28] uses Petri nets to identify some common problems in multithreaded programs such as data races, lost signals, and deadlocks. However the model introduced for condition variables in this approach only covers the communication of two threads and it is not clear how it deals with programs having more than two threads communicating through condition variables. Recently, [6, 9] have introduced an approach ensuring that every thread synchronizing under a set of condition variables eventually exits the synchronization block if that thread eventually reaches that block. This approach succeeds in verifying one of the applications of condition variables, namely the buffer. However, since this approach is not modular and relies on a Petri net analysis tool to solve the termination problem, it suffers from a long verification time when the size of the state space is increased, such that the verification of a buffer application having 20 producer and 18 consumer threads, for example, takes more than two minutes.

There are several verification techniques and type systems to check deadlock-freedom of programs that either synchronize via locks [10, 32, 47] or communicate via messages [7, 30]. Kobayashi [30, 29] proposed a type system for deadlock-free processes, ensuring that a well-typed process that is annotated with a finite capability level is deadlock-free. He extended channel types with the notion of usages, describing how often and in which order a channel is used for input and output. In his type system, which works in the context of π -calculus, the send operation is synchronous; i.e. the thread sending on a channel is suspended until the sent message is received by another thread. However, this approach and other approaches, such as [31, 5, 38], verifying deadlock-freedom in the context of π -calculus are not straight forwardly applicable to imperative programming languages.

Obligations

Inspired by the notion of capabilities [30, 29] and implicit dynamic frames [46, 44, 45], Leino *et al.* [33] later integrated deadlock prevention into a verification system for an object-oriented and imperative programming language. In this approach each thread trying to receive a message from a channel must spend one credit for that channel, where a credit for a channel is obtained if a thread is obliged to discharge an obligation for that channel. A thread can discharge an obligation for a channel if it either sends a message on that channel or delegates that obligation to another thread. This approach supports asynchronous send operations, where sending on a channel does not suspend the sender thread and there might be a state where a message is sent but not received by any thread. The notion of obligations is used in other verification approaches, which verify deadlock-freedom of semaphores [21], monitors [15, 16] and channels in a separation logic-based system [23, 24], and finite blocking in non-

terminating programs [2]. However, these approaches allow obligations to be transferred only when a thread is forked. In other words, unlike permissions which can be transferred through synchronizations, transferring obligations through synchronizations in these approaches is forbidden. In this paper we provide two mechanisms that allow these approaches to transfer obligations, along with permissions, through synchronization. Our approach can be used to verify deadlock-freedom of imperative programs where some obligation must be transferred through notifications and channels (with asynchronous send operations).

5 Conclusion

This paper introduces two techniques to transfer obligations through synchronization, while ensuring that there is no state where the transferred obligations are lost, i.e. where they are discharged from the sender thread and not loaded onto the receiver thread yet. These techniques allow the obligation-based verification approaches, which modularly verify deadlock-freedom and liveness properties of programs, to transfer obligations, along with permissions, between threads, enabling them to verify a wider range of interesting programs, where obligations must be transferred through synchronizations. We encoded the proposed proof rules in the VeriFast program verifier and succeeded in verifying deadlock-freedom of a number of interesting programs, such as some variations of client-server programs, fair readers-writers locks and dining philosophers, which cannot be modularly verified without such transfer. Integrating the two mechanisms introduced in this paper is an area of future work. Additionally, designing a new variant of Iris for programs with primitive blocking constructs on top of which our presented approach can be built is another important area of future work.

References

- 1 Aleš Bizjak, Daniel Gratzer, Robbert Krebbers, and Lars Birkedal. Iron: Managing obligations in higher-order concurrent separation logic. *To appear in POPL 2019: ACM SIGPLAN Symposium on Principles of Programming Languages, Lissabon, Portugal*, 2019.
- 2 Pontus Boström and Peter Müller. *Modular verification of finite blocking in non-terminating programs*, volume 37. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2015.
- 3 John Boyland. Checking interference with fractional permissions. In *International Static Analysis Symposium*, pages 55–72. Springer, 2003.
- 4 Pedro da Rocha Pinto, Thomas Dinsdale-Young, Philippa Gardner, and Julian Sutherland. Modular termination verification for non-blocking concurrency. In *ESOP*, pages 176–201, 2016.
- 5 Ornela Dardha and Simon J Gay. A new linear logic for deadlock-free session-typed processes. In *International Conference on Foundations of Software Science and Computation Structures*, pages 91–109. Springer, 2018.
- 6 Pedro de Carvalho Gomes, Dilian Gurov, and Marieke Huisman. Specification and verification of synchronization with condition variables. In *International Workshop on Formal Techniques for Safety-Critical Systems*, pages 3–19. Springer, 2016.
- 7 Mariangiola Dezani-Ciancaglini, Dimitris Mostrous, Nobuko Yoshida, and Sophia Drossopoulou. Session types for object-oriented languages. In *European Conference on Object-Oriented Programming*, pages 328–352. Springer, 2006.
- 8 Xinyu Feng. Local rely-guarantee reasoning. In *ACM SIGPLAN Notices*, volume 44, pages 315–327. ACM, 2009.
- 9 Pedro de C Gomes, Dilian Gurov, Marieke Huisman, and Cyrille Artho. Specification and verification of synchronization with condition variables. *Science of Computer Programming*, 163:174–189, 2018.

- 10 Colin S Gordon, Michael D Ernst, and Dan Grossman. Static lock capabilities for deadlock freedom. In *Proceedings of the 8th ACM SIGPLAN workshop on Types in language design and implementation*, pages 67–78. ACM, 2012.
- 11 Alexey Gotsman, Josh Berdine, Byron Cook, Noam Rinetzky, and Mooly Sagiv. Local reasoning for storable locks and threads. In *Asian Symposium on Programming Languages And Systems*, pages 19–37. Springer, 2007.
- 12 Christian Haack, Marieke Huisman, and Clément Hurlin. Reasoning about java’s reentrant locks. In *Asian Symposium on Programming Languages And Systems*, pages 171–187. Springer, 2008.
- 13 Christian Haack and Clément Hurlin. Separation logic contracts for a java-like language with fork/join. In *International Conference on Algebraic Methodology and Software Technology*, pages 199–215. Springer, 2008.
- 14 Jafar Hamin and Bart Jacobs. Modular verification of termination and execution time bounds using separation logic. In *Information Reuse and Integration (IRI), 2016 IEEE 17th International Conference on*, pages 110–117. IEEE, 2016.
- 15 Jafar Hamin and Bart Jacobs. Deadlock-free monitors. In *European Symposium on Programming*, pages 415–441. Springer, 2018.
- 16 Jafar Hamin and Bart Jacobs. Deadlock-free monitors: extended version. *TR CW712, Department of Computer Science, KU Leuven, Belgium. Full version available at <http://www.cs.kuleuven.be/publicaties/rapporten/cw/CW712.abs.html>*, 2018.
- 17 Jafar Hamin and Bart Jacobs. Deadlock-free monitors and channels. zenodo, <http://doi.org/10.5281/zenodo.3241454>, 2019. URL: <http://doi.org/10.5281/zenodo.3241454>, doi:10.5281/zenodo.3241454.
- 18 Tony Hoare and Peter O’Hearn. Separation logic semantics for communicating processes. *Electronic Notes in Theoretical Computer Science*, 212:3–25, 2008.
- 19 Aquinas Hobor, Andrew W Appel, and Francesco Zappa Nardelli. Oracle semantics for concurrent separation logic. In *European Symposium on Programming*, pages 353–367. Springer, 2008.
- 20 Jan Hoffmann, Michael Marmar, and Zhong Shao. Quantitative reasoning for proving lock-freedom. In *Logic in Computer Science (LICS), 2013 28th Annual IEEE/ACM Symposium on*, pages 124–133. IEEE, 2013.
- 21 Bart Jacobs. Provably live exception handling. In *Proceedings of the 17th Workshop on Formal Techniques for Java-like Programs*, page 7. ACM, 2015.
- 22 Bart Jacobs. Verifast 18.02. zenodo, <http://doi.org/10.5281/zenodo.1182724>, 2018. URL: <http://doi.org/10.5281/zenodo.1182724>, doi:10.5281/zenodo.1182724.
- 23 Bart Jacobs, Dragan Bosnacki, and Ruurd Kuiper. Modular termination verification. In *LIPICs-Leibniz International Proceedings in Informatics*, volume 37. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2015.
- 24 Bart Jacobs, Dragan Bosnacki, and Ruurd Kuiper. Modular termination verification of single-threaded and multithreaded programs. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 40(3):12, 2018.
- 25 Bart Jacobs, Jan Smans, Pieter Philippaerts, Frédéric Vogels, Willem Penninckx, and Frank Piessens. Verifast: A powerful, sound, predictable, fast verifier for c and java. *NASA Formal Methods*, 6617:41–55, 2011.
- 26 Bart Jacobs, Jan Smans, and Frank Piessens. A quick tour of the verifast program verifier. *Programming Languages and Systems*, pages 304–311, 2010.
- 27 Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal, and Derek Dreyer. Iris: Monoids and invariants as an orthogonal basis for concurrent reasoning. *ACM SIGPLAN Notices*, 50(1):637–650, 2015.
- 28 Krishna M Kavi, Alireza Moshtaghi, and Deng-Jyi Chen. Modeling multithreaded applications using petri nets. *International Journal of Parallel Programming*, 30(5):353–371, 2002.

- 29 Naoki Kobayashi. A type system for lock-free processes. *Information and Computation*, 177(2):122–159, 2002.
- 30 Naoki Kobayashi. A new type system for deadlock-free processes. In *International Conference on Concurrency Theory*, pages 233–247. Springer, 2006.
- 31 Naoki Kobayashi and Cosimo Laneve. Deadlock analysis of unbounded process networks. *Information and Computation*, 252:48–70, 2017.
- 32 Duy-Khanh Le, Wei-Ngan Chin, and Yong-Meng Teo. An expressive framework for verifying deadlock freedom. In *International Symposium on Automated Technology for Verification and Analysis*, pages 287–302. Springer, 2013.
- 33 K Rustan M Leino, Peter Müller, and Jan Smans. Deadlock-free channels and locks. In *European Symposium on Programming*, pages 407–426. Springer, 2010.
- 34 Hongjin Liang and Xinyu Feng. Progress of concurrent objects with partial methods. *Proceedings of the ACM on Programming Languages*, 2(POPL):20, 2017.
- 35 Hongjin Liang, Xinyu Feng, and Zhong Shao. Compositional verification of termination-preserving refinement of concurrent programs. In *Proceedings of the Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, page 65. ACM, 2014.
- 36 Aleksandar Nanevski, Ruy Ley-Wild, Ilya Sergey, and Germán Andrés Delbianco. Communicating state transition systems for fine-grained concurrent resources. In *European Symposium on Programming Languages and Systems*, pages 290–310. Springer, 2014.
- 37 Peter W O’Hearn. Resources, concurrency, and local reasoning. *Theoretical computer science*, 375(1-3):271–307, 2007.
- 38 Luca Padovani. Type-based analysis of linear communications. *Behavioural Types: from Theory to Tools*, page 193, 2017.
- 39 Corneliu Popeea and Andrey Rybalchenko. Compositional termination proofs for multi-threaded programs. In *TACAS*, volume 12, pages 237–251. Springer, 2012.
- 40 David Pym and Chris Tofts. A calculus and logic of resources and processes. *Formal Aspects of Computing*, 18(4):495–517, 2006.
- 41 Azalea Raad, Jules Villard, and Philippa Gardner. Colosl: Concurrent local subjective logic. In *European Symposium on Programming Languages and Systems*, pages 710–735. Springer, 2015.
- 42 John C Reynolds. Separation logic: A logic for shared mutable data structures. In *Logic in Computer Science, 2002. Proceedings. 17th Annual IEEE Symposium on*, pages 55–74. IEEE, 2002.
- 43 Reuben NS Rowe and James Brotherston. Automatic cyclic termination proofs for recursive procedures in separation logic. In *Proceedings of the 6th ACM SIGPLAN Conference on Certified Programs and Proofs*, pages 53–65. ACM, 2017.
- 44 Jan Smans, Bart Jacobs, and Frank Piessens. Implicit dynamic frames. In *Proceedings of the 10th ECOOP Workshop on Formal Techniques for Java-like Programs*, pages 1–12, 2008.
- 45 Jan Smans, Bart Jacobs, and Frank Piessens. Implicit dynamic frames: Combining dynamic frames and separation logic. In *European Conference on Object-Oriented Programming*, pages 148–172. Springer, 2009.
- 46 Jan Smans, Bart Jacobs, and Frank Piessens. Implicit dynamic frames. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 34(1):2, 2012.
- 47 Kohei Suenaga. Type-based deadlock-freedom verification for non-block-structured lock primitives and mutable references. In *Asian Symposium on Programming Languages and Systems*, pages 155–170. Springer, 2008.
- 48 Viktor Vafeiadis and Matthew Parkinson. A marriage of rely/guarantee and separation logic. In *International Conference on Concurrency Theory*, pages 256–271. Springer, 2007.
- 49 Chao Wang and Kevin Hoang. Precisely deciding control state reachability in concurrent traces with limited observability. In *VMCAI*, pages 376–394. Springer, 2014.

A Proof of Conditional Server Channels

In this section the verification of the program shown in Figure 11 is illustrated in Figures 20 and 21.

B Proof of Dining Philosophers

Transferring obligations through notifications allows to verify some other interesting programs such as some variants of *dining philosophers*, where a number of philosophers sit at a round table, think, get hungry, and eat if their neighbors are not eating (due to the limited number of forks). An implementation of this program is shown in Figures 22⁷, and 23, where a `dining_philosophers` structure consists of *phs*, a circular doubly linked list of `philosopher`, *size*, the size of this list, and a lock (*l*), where a `philosopher` structure consists of *pre*, a pointer to the previous philosopher, *next*, a pointer to the next philosopher, a condition variable *v*, and a state which can be Thinking, Hungry, or Eating. Calling `new_dining_philosophers(size)` creates *size* philosophers which are initially thinking. Calling `pickup(dp, i)` makes the *i*th philosopher in *dp* start eating if he/she is already hungry and none of its neighbors are eating. Calling `putdown(dp, i)` makes the *i*th philosopher in *dp* stop eating if he/she is already eating, and also makes this philosopher's neighbors eat if they are hungry and their neighbors are not eating.

The specification of the routine `new_dining_philosophers` in this figure indicates that creating a number of dining philosophers *dp* produces a permission `dp(dp, cvs)`, where *cvs* is a list of condition variables associated with philosophers whose levels are assigned arbitrarily. A philosopher can try to start eating only if the level of the CV associated with that philosopher is lower than the levels of the obligations of the running thread. When this philosopher starts eating the CV associated with his/her left and right neighbors are loaded onto the bag of the obligations of the running thread. These obligations are discharged when this philosopher stops eating.

One desired invariant in this program is that the number of obligations for a CV associated with a philosopher *ph* is greater than the number of his/her neighbors which are eating, that is $\text{st}(pre.state) + \text{st}(next.state) \leq Ot(ph.v)$, where $\text{st}(\text{Eating})=1$ and $\text{st}(\text{Thinking})=\text{st}(\text{Hungry})=0$. Since the state of a hungry philosopher *ph* waiting in a suspended thread *t* is changed to Eating in a thread *t'* where its neighbor stops eating, an obligation of *ph.v* must be loaded onto the bag of the obligations of *t'* and be transferred to *t* as *t'* notifies *t*. This requires transferring obligations through notifications which is possible using our proposed extension. Accordingly, as shown in Figure 22, this program can be verified if for any condition variable *v* of a philosopher *ph*, when *v* is notified the obligations $\{ph.pre.v, ph.next.v\}$ are transferred.

C Transferring Obligations Through Channels: Soundness Proof

In this appendix we provide a formalization and soundness proof for the approach introduced in Section 2. However, unfortunately, there are a few technical differences between this formalization and the one proposed in Section 2 such that in this formalization the ghost information, such as level and transferred permissions and obligations, are associated with

⁷ For simplicity, in the proof of this program we avoid writing the heap ownership permissions

22:28 Transferring Obligations Through Synchronizations

```

Mch' ::= λm. channel(snd(m), Mch, M'ch) * trandit(snd(m)) ∧ R(snd(m))=1 ∧
Mr(snd(m))={1}
M'ch' ::= λm. {snd(m)}

Mch ::= λm. channel(snd(m), Mch', M'ch') * (fst(m)=done ? true : trandit(snd(m))) ∧
Mr(snd(m))={1} ∧ ¬S(snd(m))
M'ch ::= λm. fst(m)=done ? {} : {snd(m)}

Ms ::= λm. channel(snd(m), Mch', M'ch') * trandit(snd(m)) ∧ Mr(snd(m))={1} ∧ ¬S(snd(m))
M's ::= λm. {snd(m)}

routine server(channel s){
req : {obs({}, {s∞}) * channel(s, Ms, M's) ∧ Mr(s)={1} ∧ S(s)}
while(true){
inv : {obs({}, {s∞})}
(thr, ch') := receive(s);
{obs({ch'}, {s∞}) * channel(ch', Mch', M'ch') * trandit(ch') ∧ Mr(ch')={1} ∧ ¬S(ch')}
ch := new_channel();
{obs({ch'}, {s∞}) * trandit(ch') * channel(ch, Mch, M'ch) ∧ R(ch)=1 ∧ Mr(ch)={1} ∧
¬S(ch)}
g_credit(ch); g_trandit(ch);
{obs({ch', ch}, {s∞, ch}) * trandit(ch') * credit(ch) * trandit(ch)}
send(ch', (through, ch));
{obs({}, {s∞, ch}) * credit(ch)}
fork(cserver(ch))
{obs({}, {s∞})} }
ens : {false}

routine client(channel s){
req : {obs({}, {}) * channel(s, Ms, M's) * trandit(s)}
ch' := new_channel();
{obs({}, {}) * trandit(s) * channel(ch', Mch', M'ch') ∧ R(ch')=1 ∧ Mr(ch')={1} ∧ ¬S(ch')}
g_credit(ch'); g_trandit(ch');
{obs({ch'}, {ch'}) * trandit(s) * credit(ch') * trandit(ch')}
send(s, (through, ch'));
{obs({}, {ch'}) * credit(ch')}
(thr, ch) := receive(ch'); if(thr≠through) abort();
{obs({ch}, {}) * channel(ch, Mch, M'ch) * trandit(ch) ∧ R(ch)=1 ∧ Mr(ch)={1}}
g_credit(ch'); g_trandit(ch');
{obs({ch, ch'}, {ch'}) * trandit(ch) * credit(ch') * trandit(ch')}
send(ch, (request(), ch'));
{obs({}, {ch'}) * credit(ch')}
(res, ch1) := receive(ch'); if(res=through ∨ ch≠ch1) abort();
{obs({ch}, {}) * trandit(ch)}
send(ch, (done, ch'))
ens : {obs({}, {})}

```

■ **Figure 20** Verification of the program in Figure 11 (part one of two)

```

routine main(){
req : {obs( $\{\}$ ,  $\{\}$ )
   $s := \text{new\_channel}()$ ;
  {obs( $\{\}$ ,  $\{\}$ ) * channel( $s, Ms, M's \wedge M^r(s)=\{1\} \wedge S(s)$ )
  g_trandits( $s$ );
  {obs( $\{\}$ ,  $\{s^\infty\}$ ) * trandit $^\infty(s)$ }
  fork(
    {obs( $\{\}$ ,  $\{s^\infty\}$ )}
    server( $s$ );
    {obs( $\{\}$ ,  $\{\}$ ) * trandit $^\infty(s)$ }
  fork(client( $s$ ));
  {obs( $\{\}$ ,  $\{\}$ ) * trandit $^\infty(s)$ }
  client( $s$ )
ens : {obs( $\{\}$ ,  $\{\}$ )}}

```

■ **Figure 21** Verification of the program in Figure 11 (part two of two)

channel addresses via the `channel` permissions rather than via global functions⁸. The proof rules associated with this formalization and the verification of the program in Figure 3, proved using these rules, are shown in Sections C.4 and C.6, respectively.

C.1 Syntax and Semantics of Programs

The syntax of our programming language is defined in Figure 24, where `val(e)` is a command that simply yields the value of e as its result and has no side effects, `let(x, c_1, c_2)` is syntactic sugar for $x:=c_1; c_2$, `fork(c)` creates a thread executing c , `while(c)` keeps executing c while c evaluates to `true`, `if(e, c_1, c_2)` executes c_1 if e evaluates to `true` and otherwise it executes c_2 , `send` and `receive` are used for sending and receiving on channels, and `wait`, which cannot be used by programmers, indicates that the related thread has tried to receive from an empty channel. Additionally, instead of defining three ghost commands `g_credit`, `g_trandit`, and `g_trandits`, we define a single ghost command `nop` which is inserted into the program for verification purposes and has no effect on the program's behavior. The expressions used in the syntax of programs can be evaluated and substituted as shown in Figure 25.

The small step semantics, defined in Figure 27, relates two *configurations*, defined in Figure 26. A configuration consists of a heap, which maps a channel identifier to the list of messages of that channel; a thread table, which maps a thread identifier to the pair command-context related to that thread; and a list of server channels.

22:30 Transferring Obligations Through Synchronizations

```

dp(dining_philosophers dp, list(waitobj) cvs) =
  lock(dp.l) * phs_cvs(dp.phs, dp.phs.pre, cvs) ∧
  I(dp.l)=linv(dp) ∧ L(dp.l)=new_dp ∧ size(cvs)=dp.size ∧ ∀0≤i≤size. R(rl) < R(rs[i])

```

```

phs_cvs(ph, ph', cvs) =
  ph=ph' ? cvs=[ph.v] : ∃cvs', cvs=[ph.v :: cvs'] * phs_cvs(ph.next, ph', cvs')

```

```

linv(dp dp) = λWt. λOt. philosophers(dp.phs, dp.phs.pre, Wt, Ot, dp.l)

```

```

philosophers(ph, ph', Wt, Ot, l) = philosopher(ph, Wt, Ot, l, ph.pre, ph.next) *
  ph=ph' ? true : philosophers(ph.next, ph', Wt, Ot, l)

```

```

philosopher(ph, Wt, Ot, l, pre, next) = cond(ph.v, true, {pre.v, next.v}) ∧
  pre=ph.pre ∧ next=ph.next ∧ ph≠pre ∧ ph≠next ∧ pre≠next ∧
  next.pre=ph ∧ pre.next=ph ∧ L(ph.v)=l ∧ 0 ≤ Wt(ph.v) ∧ Wt(ph.v) ≤ 1 ∧
  (ph.state=Hungry ∨ Wt(ph.v) ≤ 0) ∧ (ph.state≠Hungry ∨ 0 < Wt(ph.v)) ∧
  (Wt(ph.v) ≤ 0 ∨ 0 < st(pre.state) + st(next.state)) ∧
  st(pre.state) + st(next.state) ≤ Ot(ph.v)

```

```

st(state) ::= state=Eating ? 1 : 0

```

```

routine new_dining_philosophers(int size){
  req : {∀0≤i<size. rl < rs[i] ∧ size(rs)=size ∧ 2 < size}
  ens : {λdp. dp(dp, cvs) ∧ size(cvs)=size ∧ R(dp.l)=rl ∧ ∀0≤i<size. R(cvs[i])=rs[i]}
  phs := philosopher(state:=new_cell(Thinking), v:=new_cvar,
    pre:=new_cell(null), next:=new_cell(null));
  ph1 := philosopher(state:=new_cell(Thinking), v:=new_cvar,
    pre:=new_cell(null), next:=new_cell(null));
  ph2 := philosopher(state:=new_cell(Thinking), v:=new_cvar,
    pre:=new_cell(null), next:=new_cell(null));
  phs.pre := ph2; phs.next := ph1;
  ph1.pre := phs; ph1.next := ph2;
  ph2.pre := ph1; phs.next := phs;
  l := new_lock; i := 3
  while(i<size){
    ph:=philosopher(state:=new_cell(Thinking), v:=new_cvar, pre:=phs, next:=phs.next);
    phs.next.pre := ph;
    phs.next = ph;
    i := i+1
  };
  dining_philosophers(phs:=phs, l:=l, size:=size)
}

```

■ **Figure 22** Dining philosophers (part one of two)

```

routine pickup(dining_philosophers dp, int i){
  req : {obs(O) * dp(dp, cvs) ∧ cvs[i] < O ∧ 0 ≤ i < size(cvs)}
  ens : {obs(O ⊕ {cvs[(i+n-1)%n], cvs[(i+1)%n]} * dp(dp, cvs) ∧ n=size(cvs)}
  acquire(dp.l);
  ph := dp.phs;
  j := new_int(0);
  while(j < i){
    ph := ph.next;
    j := j+1
  };
  if(ph.state ≠ Thinking)
    abort;
  ph.state := Hungry;
  if(ph.next.state=Eating ∨ ph.pre.state=Eating)
    wait(ph.v, dp.l)
  else
    ph.state := Eating;
  release(dp.l)}

routine putdown(dining_philosophers dp, int i){
  req : {obs(O) * dp(dp, cvs) ∧ 0 ≤ i < size(cvs)}
  ens : {obs(O - {cvs[(i+n-1)%n], cvs[(i+1)%n]} * dp(dp, cvs) ∧ n=size(cvs)}
  acquire(dp.l);
  ph := dp.phs;
  j := new_int(0);
  while(j < i){
    ph := ph.next;
    j := j+1
  };
  if(ph.state ≠ Eating)
    abort;
  ph.state := Thinking;
  test_and_notify(ph.next);
  test_and_notify(ph.pre);
  release(dp.l)}

routine test_and_notify(philosopher ph)
  {if(ph.next.state≠Eating ∧ ph.pre.state≠Eating ∧ ph.state=Hungry){
    ph.state := Eating;
    notify(ph.v)
  }}

```

■ **Figure 23** Dining philosophers (part two of two)

$$\begin{aligned}
& c \in \text{Commands}, e \in \text{Expressions}, z \in \mathbb{Z}, b \in \text{Booleans}, x \in \text{Variables} \\
& e ::= z \mid x \mid e_1 + e_2 \mid (e_1, e_2) \mid \text{fst}(e) \mid \text{snd}(e) \\
& \quad \mid \text{true} \mid e_1 = e_2 \mid e_1 \leq e_2 \mid \neg e \\
& c ::= \text{val}(e) \mid \text{let}(x, c_1, c_2) \mid \text{fork}(c) \mid \text{while}(c) \mid \text{if}(e, c_1, c_2) \\
& \quad \mid \text{new_channel}(b) \mid \text{send}(e_1, e_2) \mid \text{receive}(e) \mid \text{wait}(e) \mid \text{nop}
\end{aligned}$$

■ **Figure 24** Syntax of the programming language

C.2 Syntax and Semantics of Assertions

The syntax of assertions is defined in Figure 28⁹. Note that the *location* of a channel ch consists of the *obligation* of ch and the permissions and the obligations which are transferred through a specific message sent on ch , denoted by $M(ch)$ and $M'(ch)$. Also note that the permissions which are transferred through a specific message sent on ch , denoted by $M(ch)$, are specified through an index (as well as the required arguments) pointing to a table in which each element is a function that given a list of arguments and a message returns an assertion. This makes it possible for the predicates specifying these permissions to recursively refer to themselves or to each other, as in Figure 20¹⁰. The obligation of a location ch , denoted by $O(ch)$, consists of the address of that location, denoted by $A(ch)$, as well as other related information such as the level of ch , denoted by $R(ch)$; the bag of the levels of the obligations which are possibly imported by ch , denoted by $M'(ch)$; and whether ch is a server channel or not, denoted by $S(ch)$.

The proposed assertions describe some ghost resources, namely $p, \tilde{O}, \tilde{I}, C, T$, that keep track of allocated channels, and the current thread's obligations, importers, credits, and trandits, respectively, shown in Figure 29.

C.3 Weakest Precondition of Commands

The weakest precondition of a command c for $n > 0$ steps w.r.t. a postcondition Q (with a given predicate table, specified by pt), denoted by $\text{wp}_{n,pt}(c, Q)$ is defined in Figure 30. Note that $\text{wp}(c, Q)_{0,pt} = \text{true}$. Also note that for the sake of simplicity the index pt is elided. Having this definition, we define the weakest precondition of a context and the weakest precondition of a command-context as shown in Definitions 1 and 2. Having these definitions, we can prove some auxiliary lemmas, shown in Lemmas 4, 5, 6, and 7, which are used to prove Theorem 13.

⁸ The reason is to make this formalization consistent with the one in Appendix D which is machine-checked with Coq, where ghost information is associated with lock and condition variable addresses via the `lock` and `cond` permissions. However, we believe one way to formalize the precise approach of Section 3 would be to define assertions as functions from ghost information to separating conjunctions of chunks. In the soundness proof, one would track these as partial functions whose domain is the set of allocated addresses. The functions passed into the assertions would be *totalizations* of these partial functions. An assertion is true if it is true for all totalizations of the functions.

⁹ Note that we use a shallow embedding: assertions have no variables; to model quantifications, we use meta-level functions from values to assertions.

¹⁰ An alternative approach is to use a step-indexed domain of assertions, as in Iris [27]. There, \blacktriangleright *Assertions* could be used instead of $\text{Indexes} \times \text{Lists}(\text{Arguments})$, where \blacktriangleright is Iris's *guard* for *guarded recursive definitions*.

$$\begin{aligned}
v &\in \text{Values} ::= z \mid b \mid (v, v) \\
\llbracket \cdot \rrbracket &\in \text{Expressions} \rightarrow \text{Values} \\
\llbracket z \rrbracket &= z \\
\llbracket x \rrbracket &= 0 \\
\llbracket e_1 + e_2 \rrbracket &= \llbracket e_1 \rrbracket + \llbracket e_2 \rrbracket \\
\llbracket (e_1, e_2) \rrbracket &= (\llbracket e_1 \rrbracket, \llbracket e_2 \rrbracket) \\
\llbracket \text{true} \rrbracket &= \text{true} \\
\llbracket (e_1 = e_2) \rrbracket &= (\llbracket e_1 \rrbracket = \llbracket e_2 \rrbracket) \\
\llbracket (e_1 \leq e_2) \rrbracket &= (\llbracket e_1 \rrbracket \leq \llbracket e_2 \rrbracket) \\
\llbracket (\neg e) \rrbracket &= (\neg \llbracket e \rrbracket) \\
\llbracket \text{fst}(e) \rrbracket &= \begin{cases} \llbracket e_1 \rrbracket & \text{if } e = (e_1, e_2) \\ 0 & \text{otherwise} \end{cases} \\
\llbracket \text{snd}(e) \rrbracket &= \begin{cases} \llbracket e_2 \rrbracket & \text{if } e = (e_1, e_2) \\ 0 & \text{otherwise} \end{cases} \\
z[v/x] &= z \\
x[v/x'] &= x \quad \text{if } x \neq x' \\
x[v/x] &= v \\
(e_1 + e_2)[v/x] &= e_1[v/x] + e_2[v/x] \\
(e_1, e_2)[v/x] &= (e_1[v/x], e_2[v/x]) \\
\text{true}[v/x] &= \text{true} \\
(e_1 = e_2)[v/x] &= (e_1[v/x] = e_2[v/x]) \\
(e_1 \leq e_2)[v/x] &= (e_1[v/x] \leq e_2[v/x]) \\
(\neg e)[v/x] &= (\neg e[v/x]) \\
\text{fst}(e)[v/x] &= \text{fst}(e[v/x]) \\
\text{snd}(e)[v/x] &= \text{snd}(e[v/x]) \\
\text{val}(e)[v/x] &= \text{val}(e[v/x]) \\
\text{let}(x, c_1, c_2)[v/x] &= \text{let}(x, c_1, c_2) \\
\text{let}(x, c_1, c_2)[v/x'] &= \text{let}(x, c_1[v/x'], c_2[v/x']) \quad \text{if } x \neq x' \\
\text{fork}(c)[v/x] &= \text{fork}(c[v/x]) \\
\text{while}(c)[v/x] &= \text{while}(c[v/x]) \\
\text{if}(e, c_1, c_2)[v/x] &= \text{if}(e[v/x], c_1[v/x], c_2[v/x]) \\
\text{new_channel}(b)[v/x] &= \text{new_channel}(b) \\
\text{send}(e_1, e_2)[v/x] &= \text{send}(e_1[v/x], e_2[v/x]) \\
\text{receive}(e)[v/x] &= \text{receive}(e[v/x]) \\
\text{nop}[v/x] &= \text{nop}
\end{aligned}$$

■ **Figure 25** Evaluation of expressions and substitution of expressions and commands

22:34 Transferring Obligations Through Synchronizations

$$\begin{aligned}
m &\in \text{Messages} = \text{Values} \\
\text{Addresses} &= \mathbb{Z} \\
\text{ThreadId}s &= \mathbb{Z} \\
h &\in \text{Heaps} = \text{Addresses} \rightarrow \text{Lists}(\text{Messages}) \\
\xi &\in \text{Contexts} ::= \text{done} \mid \text{let}'(x, c, \xi) \mid \text{while}'(c, \xi) \\
\theta &\in \text{ThreadConfigurations} ::= (c; \xi) \\
t &\in \text{ThreadTables} = \text{ThreadId}s \rightarrow \text{ThreadConfigurations} \\
s &\in \text{ServerChannelsSets} = \text{Sets}(\text{Addresses}) \\
\kappa &\in \text{Configurations} = \text{Heaps} \times \text{ThreadTables} \times \text{ServerChannelsSets}
\end{aligned}$$

■ **Figure 26** Configurations

$$\begin{aligned}
(t[id:=\text{new_channel}(b); \xi], h[z:=\emptyset], s) &\rightsquigarrow (t[id:=\text{val}(z); \xi], h[z:=[]], b=\text{true} ? s \cup \{z\} : s) \\
(t[id:=\text{send}(e_1, e_2); \xi], h[[e_1]:=M], s) &\rightsquigarrow (t[id:=\text{tt}; \xi], h[[e_1]:=M \cdot [[e_2]]], s) \\
(t[id:=\text{receive}(e); \xi], h[[e]:=[m] \cdot M], s) &\rightsquigarrow (t[id:=\text{val}(m); \xi], h[[e]:=M], s) \\
(t[id:=\text{receive}(e); \xi], h[[e]:=[]], s) &\rightsquigarrow (t[id:=\text{wait}(e); \xi], h[[e]:=[]], s) \\
(t[id:=\text{wait}(e); \xi], h[[e]:=[m] \cdot M], s) &\rightsquigarrow (t[id:=\text{val}(m); \xi], h[[e]:=M], s) \\
(t[id:=\text{fork}(c); \xi, id':=\emptyset], h, s) &\rightsquigarrow (t[id:=\text{tt}; \xi, id':=c; \text{done}], h, s) \\
(t[id:=\text{let}(x, c_1, c_2); \xi], h, s) &\rightsquigarrow (t[id:=c_1; \text{let}'(x, c_2, \xi)], h, s) \\
(t[id:=\text{val}(e); \text{let}'(x, c, \xi)], h, s) &\rightsquigarrow (t[id:=c[[e]]/x; \xi], h, s) \\
(t[id:=\text{val}(e); \text{done}], h, s) &\rightsquigarrow (t[id:=\emptyset], h, s) \\
(t[id:=\text{if}(e, c_1, c_2); \xi], h, s) &\rightsquigarrow (t[id:=c_1; \xi], h, s) \quad \text{if } [[e]] = \text{true} \\
(t[id:=\text{if}(e, c_1, c_2); \xi], h, s) &\rightsquigarrow (t[id:=c_2; \xi], h, s) \quad \text{if } [[e]] \neq \text{true} \\
(t[id:=\text{while}(c); \xi], h, s) &\rightsquigarrow (t[id:=c; \text{while}'(c, \xi)], h, s) \\
(t[id:=\text{val}(e); \text{while}'(c, \xi)], h, s) &\rightsquigarrow (t[id:=c; \text{while}'(c, \xi)], h, s) \quad \text{if } [[e]] = \text{true} \\
(t[id:=\text{val}(e); \text{while}'(c, \xi)], h, s) &\rightsquigarrow (t[id:=\text{tt}; \xi], h, s) \quad \text{if } [[e]] \neq \text{true} \\
(t[id:=\text{nop}; \xi], h, s) &\rightsquigarrow (t[id:=\text{tt}; \xi], h, s)
\end{aligned}$$

■ **Figure 27** Semantics of programs, where tt stands for $\text{val}(0)$, and $[m]$ represents a list with one element m , and $M_1 \cdot M_2$ appends two lists M_1 and M_2 .

$$\begin{aligned}
n &\in \mathbb{N} \\
\mathbb{NF} &::= n \mid \infty \\
\text{Bags}(A) &= A \rightarrow \mathbb{NF} \\
\text{Indexes} &= \mathbb{Z} \\
\text{Arguments} &= \mathbb{Z} \\
r &\in \text{Levels} = \mathbb{R} \\
o &\in \text{Obligations} = \text{Addresses} \times \text{Levels} \times \text{Bags}(\text{Levels}) \times \text{Booleans} \\
l &\in \text{Locations} = \\
&\quad \text{Obligations} \times (\text{Indexes} \times \text{Lists}(\text{Arguments})), \text{Messages} \rightarrow \text{Bags}(\text{Obligations}) \\
O, I &\in \text{Bags}(\text{Obligations}) \\
b &\in \text{Booleans} \\
\hat{v} &\in \text{AValues} ::= z \mid r \mid b \mid l \mid o \mid O \\
\alpha &\in \text{AValues} \rightarrow \text{Assertions} \\
a &\in \text{Assertions} ::= \text{channel}(l) \mid \text{credit}(z) \mid \text{trandit}(z) \mid \text{trandit}^\infty(z) \mid \text{obs}(O, I) \\
&\quad \mid b \mid a_1 \wedge a_2 \mid a_1 \vee a_2 \mid a_1 * a_2 \mid a_1 \multimap a_2 \mid \forall \alpha \mid \exists \alpha \\
\text{pt} &: \text{PredicateTables} = \text{Indexes} \rightarrow \text{Lists}(\text{Arguments}) \rightarrow \text{Messages} \rightarrow \text{Assertions} \\
O &: \text{Locations} \rightarrow \text{Obligations}, \text{ where } O((A, R, M^r, S), M, M') = (A, R, M^r, S) \\
A &: \text{Locations} \rightarrow \text{Addresses}, \text{ where } A((A, R, M^r, S), M, M') = A \\
R &: \text{Locations} \rightarrow \text{Levels}, \text{ where } R((A, R, M^r, S), M, M') = R \\
M^r &: \text{Locations} \rightarrow \text{Bags}(\text{Levels}), \text{ where } M^r((A, R, M^r, S), M, M') = M^r \\
S &: \text{Locations} \rightarrow \text{Booleans}, \text{ where } S((A, R, M^r, S), M, M') = S \\
M &: \text{Locations} \rightarrow \text{Messages} \rightarrow \text{Assertions} \\
&\quad \text{where } M((A, R, M^r, S), (M_1, M_2), M') = \text{pt}(M_1, M_2) \\
M' &: \text{Locations} \rightarrow \text{Messages} \rightarrow \text{Bags}(\text{Obligations}), \text{ where } M'((A, R, M^r, S), M, M') = M' \\
\text{Ro} &: \text{Obligations} \rightarrow \text{Levels}, \text{ where } \text{Ro}(A, R, M^r, S) = R
\end{aligned}$$

■ **Figure 28** Syntax of assertions

$p \in \text{PermissionHeaps} = \text{Locations} \rightarrow \{\text{channel}\}$
 $\text{Option}(A) ::= s \mid \emptyset$, where $s \in A$
 $\tilde{O}, \tilde{I} \in \text{Option}(\text{Bags}(\text{Obligations}))$
 $C, T \in \text{Bags}(\text{Addresses})$

$p, \tilde{O}, \tilde{I}, C, T \models \text{channel}(l) \Leftrightarrow p(l) = \text{channel}$
 $p, \tilde{O}, \tilde{I}, C, T \models \text{credit}(z) \Leftrightarrow 0 < C(z)$
 $p, \tilde{O}, \tilde{I}, C, T \models \text{trandit}(z) \Leftrightarrow 0 < T(z)$
 $p, \tilde{O}, \tilde{I}, C, T \models \text{trandit}^\infty(z) \Leftrightarrow T(z) = \infty$
 $p, \tilde{O}, \tilde{I}, C, T \models \text{obs}(O, I) \Leftrightarrow \tilde{O} = O \wedge \tilde{I} = I$
 $p, \tilde{O}, \tilde{I}, C, T \models b \Leftrightarrow b = \text{true}$
 $p, \tilde{O}, \tilde{I}, C, T \models a_1 \wedge a_2 \Leftrightarrow p, \tilde{O}, \tilde{I}, C, T \models a_1 \wedge p, \tilde{O}, \tilde{I}, C, T \models a_2$
 $p, \tilde{O}, \tilde{I}, C, T \models a_1 \vee a_2 \Leftrightarrow p, \tilde{O}, \tilde{I}, C, T \models a_1 \vee p, \tilde{O}, \tilde{I}, C, T \models a_2$
 $p, \tilde{O}, \tilde{I}, C, T \models a_1 * a_2 \Leftrightarrow \exists p_1, p_2, \tilde{O}_1, \tilde{O}_2, \tilde{I}_1, \tilde{I}_2, C_1, C_2, T_1, T_2.$
 $\quad p = p_1 \uplus p_2 \wedge \tilde{O} = \tilde{O}_1 \uplus \tilde{O}_2 \wedge \tilde{I} = \tilde{I}_1 \uplus \tilde{I}_2 \wedge C = C_1 \uplus C_2 \wedge T = T_1 \uplus T_2 \wedge$
 $\quad p_1, \tilde{O}_1, \tilde{I}_1, C_1, T_1 \models a_1 \wedge p_2, \tilde{O}_2, \tilde{I}_2, C_2, T_2 \models a_2$
 $p, \tilde{O}, \tilde{I}, C, T \models a_1 * a_2 \Leftrightarrow \forall p_1, \tilde{O}_1, \tilde{I}_1, C_1, T_1. p_1, \tilde{O}_1, \tilde{I}_1, C_1, T_1 \models a_1 \wedge$
 $\quad \tilde{O} \perp \tilde{O}_1 \wedge \tilde{I} \perp \tilde{I}_1 \Rightarrow (p \uplus p_1), (\tilde{O} \uplus \tilde{O}_1), (\tilde{I} \uplus \tilde{I}_1), (C \uplus C_1), (T \uplus T_1) \models a_2$
 $p, \tilde{O}, \tilde{I}, C, T \models \forall \alpha \Leftrightarrow \forall \hat{v} \in A \text{Values}. p, \tilde{O}, \tilde{I}, C, T \models \alpha(\hat{v})$
 $p, \tilde{O}, \tilde{I}, C, T \models \exists \alpha \Leftrightarrow \exists \hat{v} \in A \text{Values}. p, \tilde{O}, \tilde{I}, C, T \models \alpha(\hat{v})$

$a_1 \vdash a_2 \Leftrightarrow (\forall p, \tilde{O}, \tilde{I}, C, T. p, \tilde{O}, \tilde{I}, C, T \models a_1 \Rightarrow p, \tilde{O}, \tilde{I}, C, T \models a_2)$

where for any $p_1, p_2 \in A \rightarrow B$ and $\tilde{O}_1, \tilde{O}_2 \in \text{Option}(\text{Bags}(A))$ and $B_1, B_2 \in \text{Bags}(A)$

$\tilde{O}_1 \perp \tilde{O}_2 \Leftrightarrow \tilde{O}_1 = \emptyset \vee \tilde{O}_2 = \emptyset$

$p_1 \uplus p_2 = \lambda v. \begin{cases} p_1(v) & \text{if } p_2(v) = \emptyset \\ p_2(v) & \text{otherwise} \end{cases} \quad \tilde{O}_1 \uplus \tilde{O}_2 = \begin{cases} \tilde{O}_1 & \text{if } \tilde{O}_2 = \emptyset \\ \tilde{O}_2 & \text{if } \tilde{O}_1 = \emptyset \\ \text{undefined} & \text{otherwise} \end{cases}$

$B_1 \uplus B_2 = \lambda v. \begin{cases} \infty & \text{if } B_1(v) = \infty \text{ or } B_2(v) = \infty \\ B_1(v) + B_2(v) & \text{otherwise} \end{cases}$

■ **Figure 29** Satisfaction relation

$$\begin{aligned}
\text{wp} &\in \text{WeakestPreconditions} = \\
&\text{Commands} \rightarrow (\text{Values} \rightarrow \text{Assertions}) \rightarrow \mathbb{N} \rightarrow \text{PredicateTables} \rightarrow \text{Assertions} \\
\text{levels}(O) &= \{\text{Ro}(o) \mid o \in O\} \\
o \prec' R &\Leftrightarrow \forall r \in R. \text{Ro}(o) < r \\
o \prec O &\Leftrightarrow o \prec' \text{levels}(O) \\
o \prec^r I &\Leftrightarrow \forall o' \in I. o = o' \vee o \prec' \text{M}^r(o') \\
\text{wp}_n(\text{val}(e), Q) &= Q(\llbracket e \rrbracket) \\
\text{wp}_n(\text{new_channel}(b), Q) &= \forall z. \exists r, M^r, M, M'. \text{channel}((z, r, M^r, b), M, M') \multimap Q(z) \\
\text{wp}_n(\text{send}(e_1, e_2)) &= \exists O, I, ch, m. (\text{obs}(O, I) * \text{channel}(ch) * \text{M}(ch)(m) * \\
&\quad (\text{M}^r(ch) = \{\!\!\! \{\} \} \vee \text{trandit}(\text{A}(ch))) \wedge \text{levels}(\text{M}^r(ch)(m)) \subseteq \text{M}^r(ch) \wedge \text{A}(ch) = \llbracket e_1 \rrbracket \wedge m = \llbracket e_2 \rrbracket) \\
&\quad * ((\text{obs}(O - \{\!\!\! \{\text{O}(ch)\}\!\!\!\}, \text{M}^r(ch)(m), I) * \text{channel}(ch)) \multimap Q(\text{tt})) \\
\text{wp}_n(\text{receive}(e)) &= \exists O, I, ch. (\text{obs}(O, I) * \text{channel}(ch) * ((\text{S}(ch) \vee \text{credit}(\text{A}(ch))) \wedge \\
&\quad \text{O}(ch) \prec O \wedge \text{O}(ch) \prec^r I \wedge (\neg \text{S}(ch) \vee (O = \{\!\!\! \{\} \} \wedge \forall o \in I. o = \text{O}(ch))) \wedge \text{A}(ch) = \llbracket e \rrbracket)) * \\
&\quad \forall m. ((\text{obs}(O \uplus \text{M}^r(ch)(m), I - \{\!\!\! \{\text{O}(ch)\}\!\!\!\}) * \text{M}(ch)(m)) \multimap Q(m)) \\
\text{wp}_n(\text{wait}(e)) &= \text{wp}_n(\text{receive}(e)) \\
\text{wp}_n(\text{fork}(c), Q) &= \exists O_1, O_2, I_1, I_2. \text{obs}(O_1 \uplus O_2, I_1 \uplus I_2) * (\text{obs}(O_1, I_1) \multimap Q(\text{tt})) * \\
&\quad (\text{obs}(O_2, I_2) \multimap \text{wp}_{n-1}(c, \lambda _. \text{obs}(\{\!\!\! \{\} \}, \{\!\!\! \{\} \}))) \\
\text{wp}_n(\text{if}(e, c_1, c_2), Q) &= (\llbracket e \rrbracket = \text{true}) ? \text{wp}_{n-1}(c_1, Q) : \text{wp}_{n-1}(c_2, Q) \\
\text{wp}_n(\text{while}(c), Q) &= \text{wp}_{n-1}(c, (\lambda vl. vl \neq \text{true} ? Q(\text{tt}) : \text{wp}_{n-1}(\text{while}(c), Q))) \\
\text{wp}_n(\text{nop}) &= \\
&\boxed{\text{as g_credit}} \\
&\quad (\exists O, I, ch. \text{obs}(O, I) * \text{channel}(ch) * \\
&\quad ((\text{obs}(O \uplus \{\!\!\! \{\text{O}(ch)\}\!\!\!\}, I) * \text{channel}(ch) * \text{credit}(\text{A}(ch))) \multimap Q(\text{tt}))) \vee \\
&\boxed{\text{as g_trandit}} \\
&\quad (\exists O, I, ch. \text{obs}(O, I) * \text{channel}(ch) * \\
&\quad ((\text{obs}(O, I \uplus \{\!\!\! \{\text{O}(ch)\}\!\!\!\}) * \text{channel}(ch) * \text{trandit}(\text{A}(ch))) \multimap Q(\text{tt}))) \vee \\
&\boxed{\text{as g_trandits}} \\
&\quad (\exists O, I, ch. \text{obs}(O, I) * \text{channel}(ch) * \\
&\quad ((\text{obs}(O, I \uplus (\lambda o. o = \text{O}(ch) ? \infty : 0)) * \text{channel}(ch) * \text{trandit}^\infty(\text{A}(ch))) \multimap Q(\text{tt})))
\end{aligned}$$

■ **Figure 30** Weakest precondition, where tt stands for 0.

22:38 Transferring Obligations Through Synchronizations

► **Definition 1** (Weakest Precondition of a Context).

$$\text{wp}_n(\xi) = \begin{cases} \lambda _. \text{obs}(\{\}, \{\}) & \text{if } \xi = \text{done} \\ \lambda v. \text{wp}_n(c[v/x], \text{wp}_n(\xi')) & \text{if } \xi = \text{let}'(x, c, \xi') \\ \lambda v. v \neq \text{true} ? \text{wp}_n(\text{tt}, \text{wp}_n(\xi')) : \text{wp}_n(c, \text{wp}_{n-1}(\xi)) & \text{if } \xi = \text{while}'(c, \xi') \end{cases}$$

► **Definition 2** (Weakest Precondition of a command-context).

$$\text{wpcx}_n(c, \xi) = \text{wp}_n(c, \text{wp}_n(\xi))$$

► **Lemma 3** (Weakening Postcondition).

$$\forall n, c, Q, Q', p, \tilde{O}, \tilde{I}, C, T. \quad p, \tilde{O}, \tilde{I}, C, T \models \text{wp}_n(c, Q) \wedge (\forall z. Q(z) \vdash Q'(z)) \Rightarrow \\ \forall n' \leq n. p, \tilde{O}, \tilde{I}, C, T \models \text{wp}_{n'}(c, Q')$$

Proof. By induction on n and case analysis of c . ◀

► **Lemma 4** (Weakest Precondition of new_channel).

$$\forall b, \xi, p, O, I, C, T. \\ p, O, I, C, T \models \text{wpcx}_n(\text{new_channel}(b), \xi) \Rightarrow \forall z. (\forall l. A(l) = z \Rightarrow p(l) = \emptyset) \Rightarrow \\ \exists r, M^r, M, M', ch. p[ch := \text{channel}], O, I, C, T \models \text{wpcx}_n(\text{val}(z), \xi) \wedge \\ A(ch) = z \wedge R(ch) = r \wedge M^r(ch) = M^r \wedge M(ch) = M \wedge M'(ch) = M' \wedge S(ch) = b$$

► **Lemma 5** (Weakest Precondition of send).

$$\forall n, e_1, e_2, \xi, p, O, I, C, T. \quad p, O, I, C, T \models \text{wpcx}_n(\text{send}(e_1, e_2), \xi) \Rightarrow \\ \exists p_1, p_2, C_1, C_2, T_1, T_2, T_e, ch, m. p = p_1 \uplus p_2 \wedge C = C_1 \uplus C_2 \wedge T = T_1 \uplus T_2 \wedge T_e = T_2(\llbracket e_1 \rrbracket) \\ \wedge A(ch) = \llbracket e_1 \rrbracket \wedge m = \llbracket e_2 \rrbracket \wedge (M^r(ch) = \{\} \vee 0 < T_e) \wedge \text{levels}(M'(ch)(m)) \subseteq M^r(ch) \wedge \\ p_2(ch) = \text{channel} \wedge \\ p_1, \emptyset, \emptyset, C_1, T_1 \models M(ch)(m) \wedge \\ (M^r(ch) = \{\} \Rightarrow p_2, O - \{\text{O}(ch)\} - M'(ch)(m), I, C_2, T_2 \models \text{wpcx}_n(\text{tt}, \xi)) \wedge \\ (M^r(ch) \neq \{\} \Rightarrow p_2, O - \{\text{O}(ch)\} - M'(ch)(m), I, C_2, T_2[\llbracket e_1 \rrbracket := T_e - 1] \models \text{wpcx}_n(\text{tt}, \xi))$$

► **Lemma 6** (Weakest Precondition of receive).

$$\forall n, e, \xi, p, O, I, C, T. \quad p, O, I, C, T \models \text{wpcx}_n(\text{receive}(e), \xi) \Rightarrow \\ \exists ch. p(ch) = \text{channel} \wedge (S(ch) = \text{true} \vee 0 < C(\llbracket e \rrbracket)) \wedge \text{O}(ch) \prec O \wedge \text{O}(ch) \prec^r I \wedge \\ (\neg S(ch) \vee (O = \{\} \wedge \forall o \in I. o = \text{O}(ch))) \wedge A(e) = \llbracket ch \rrbracket \wedge \\ \forall m, p_1, C_1, T_1. p_1, \emptyset, \emptyset, C_1, T_1 \models M(ch)(m) \Rightarrow \\ p \uplus p_1, O \uplus M'(ch)(m), I - \{\text{O}(ch)\}, C[\llbracket e \rrbracket := C(\llbracket e \rrbracket) - 1] \uplus C_1, T \uplus T_1 \models \text{wpcx}_n(\text{val}(m), \xi)$$

► **Lemma 7** (Weakest Precondition of fork).

$$\forall n, c, \xi, p, O, I, C, T. \quad p, O, I, C, T \models \text{wpcx}_n(\text{fork}(c), \xi) \Rightarrow \\ \exists p_1, p_2, O_1, O_2, I_1, I_2, C_1, C_2, T_1, T_2. \\ p = p_1 \uplus p_2 \wedge O = O_1 \uplus O_2 \wedge I = I_1 \uplus I_2 \wedge C = C_1 \uplus C_2 \wedge T = T_1 \uplus T_2 \wedge \\ p_1, O_1, I_1, C_1, T_1 \models \text{wpcx}_n(\text{tt}, \xi) \wedge \\ p_2, O_2, I_2, C_2, T_2 \models \text{wp}_{n-1}(c, \lambda _. \text{obs}(\{\}, \{\}))$$

$$\begin{array}{l}
\text{NEWCHANNEL} \\
\{\text{true}\} \text{new_channel } \{\lambda a. \text{channel}((a, r, R, b), (M_{index}, M_{args}), M')\} \\
\\
\text{SEND} \\
\{\text{obs}(O, I) * \text{channel}(ch) * M(ch)(m) * (M'(ch) = \{\}) \vee \text{trandit}(a)) \wedge \\
\text{levels}(M'(ch)(m)) \subseteq M'(ch) \wedge A(ch) = a\} \text{send}(a, m) \\
\{\lambda _. \text{obs}(O - \{\text{O}(ch)\}) - M'(ch)(m), I) * \text{channel}(ch)\} \\
\\
\text{RECEIVE} \\
\{\text{obs}(O, I) * \text{channel}(ch) * (S(ch) \vee \text{credit}(a)) \wedge \text{O}(ch) \prec O \wedge \text{O}(ch) \prec^r I \wedge \\
(\neg S(ch) \vee (O = \{\} \wedge \forall o \in I. o = \text{O}(ch))) \wedge A(ch) = a\} \text{receive}(a) \\
\{\lambda m. \text{obs}(O \uplus M'(ch)(m), I - \{\text{O}(ch)\}) * \text{channel}(ch) * M(ch)(m)\} \\
\\
\text{CREDIT} \\
\{\text{obs}(O) * \text{channel}(ch)\} \text{nop } \{\lambda _. \text{obs}(O \uplus \{\text{O}(ch)\}) * \text{channel}(ch) * \text{credit}(A(ch))\} \\
\\
\text{TRANDIT} \\
\{\text{obs}(O, I) * \text{channel}(ch)\} \text{nop } \{\lambda _. \text{obs}(O, I \uplus \{\text{O}(ch)\}) * \text{channel}(ch) * \text{trandit}(A(ch))\} \\
\\
\text{TRANDITS} \\
\{\text{obs}(O, I) * \text{channel}(ch)\} \text{nop } \{\lambda _. \text{obs}(O, I \uplus \{\text{O}(ch)^\infty\}) * \text{channel}(ch) * \text{trandit}^\infty(A(ch))\}
\end{array}$$

■ **Figure 31** The proof rules ensuring deadlock-freedom of importer channels, where ghost information is associated with channel addresses via the `channel` permissions.

C.4 Correctness of Commands

We define *correctness of commands*, as shown in Definition 8, ensuring that each proposed proof rule, where $\text{correct}_{pt}(P, c, Q)$ is abbreviated as $\{P\} c \{Q\}$, respects the definition of the weakest precondition. Having this definition we prove the proposed proof rules, ensuring deadlock freedom of importer channels, as well as some other necessary proof rules shown in Theorems 9, 10, and 11.

► **Definition 8** (Correctness of Commands). *A command is correct w.r.t a precondition P and a postcondition Q if and only if P implies the weakest precondition of that command w.r.t Q .*

$$\text{correct}_{pt}(P, c, Q) \Leftrightarrow \forall n. P \Rightarrow \text{wp}_{n,pt}(c, Q)$$

► **Theorem 9** (Rule Sequential Composition).

$$\text{correct}(P, c_1, Q) \wedge (\forall z. \text{correct}(Q(z), c_2[z/x], R)) \Rightarrow \text{correct}(P, \text{let}(x, c_1, c_2), R)$$

► **Theorem 10** (Rule Consequence).

$$\text{correct}(P, c, Q) \wedge (P' \vdash P) \wedge (\forall z. Q(z) \vdash Q'(z)) \Rightarrow \text{correct}(P', c, Q')$$

► **Theorem 11** (Rule Frame).

$$\text{correct}(P, c, Q) \Rightarrow \text{correct}(P * F, c, \lambda z. Q(z) * F)$$

As previously mentioned, since in this formalization ghost information is associated with channel addresses via the `channel` permissions rather than via global functions, we provide a new version of the proof rules, proposed in Section 2, associated with this formalization as shown in Figure 31.

C.5 Validity of a Configuration

We define *validity of a configuration*, shown in Definition 12, and prove that 1) starting from a valid configuration, all the subsequent configurations of the execution are also valid (Theorem 13), 2) a valid configuration is not deadlocked (Theorem 14), and 3) if a program c is verified by the proposed proof rules, where the verification starts from empty bags of obligations and importers and ends with such bags too, then the initial configuration, where the heap is empty, denoted by $\mathbf{0} = \lambda _ . \emptyset$, and there is only one thread with the command c (and a context done), and the list of server channels is empty, is a valid configuration (Theorem 16).

► **Definition 12** (Validity of a Configuration). *A configuration (t, h, s) is valid for n steps, denoted by $\text{valid}_n(t, h, s)$, if there exists a set of augmented threads A , consisting of the identifier (id), the program (c), the context (ξ), the permission heap (p), the obligations (O), the importers (I), the credits (C), and the trandits (T) associated with each thread such that all of the following conditions hold:*

1. $\forall id, c, \xi. t(id) = (c; \xi) \Leftrightarrow \exists p, O, I, C, T. (id, c, \xi, p, O, I, C, T) \in A$
2. $\forall (id_1, c_1, \xi_1, p_1, O_1, I_1, C_1, T_1) \in A, (id_2, c_2, \xi_2, p_2, O_2, I_2, C_2, T_2) \in A. id_1 = id_2 \Rightarrow (id_1, c_1, \xi_1, p_1, O_1, I_1, C_1, T_1) = (id_2, c_2, \xi_2, p_2, O_2, I_2, C_2, T_2)$
3. $\forall l_1, l_2. \text{Pt}(l_1) \neq \emptyset \wedge \text{Pt}(l_2) \neq \emptyset \wedge \mathbf{A}(l_1) = \mathbf{A}(l_2) \Rightarrow l_1 = l_2$
4. $\forall l. \text{Pt}(l) \neq \emptyset \Rightarrow h(\mathbf{A}(l)) \neq \emptyset$ and $\forall z. (\forall l. \mathbf{A}(l) = z \Rightarrow \text{Pt}(l) = \emptyset) \Rightarrow h(z) = \emptyset$
5. $\forall ch. \text{Pt}(ch) = \text{channel} \Rightarrow$
 - a. $\mathbf{M}^r(ch) \neq \{\} \Rightarrow \text{Tt}(\mathbf{A}(ch)) + \text{size}_h(ch) \leq \text{lt}(ch) \wedge$
 - b. $\forall m \in \text{queue}_h(ch). \text{levels}(\mathbf{M}^r(ch)(m)) \subseteq \mathbf{M}^r(ch) \wedge$
 - c. $\text{S}(ch) = \text{false} \Rightarrow$
 $\text{Ct}(\mathbf{A}(ch)) \leq \text{Ot}(ch) + \text{size}_h(ch) + \sum_{ch' \text{ where } \text{Pt}(ch') = \text{channel}} \sum_{m \in \text{queue}_h(ch')} \mathbf{M}^r(ch')(m)(ch)$
6. $\forall (id, c, \xi, p, O, I, C, T) \in A. p, O, I, C, T \models \text{wpcx}_n(c, \xi)$
7. $\forall z \in s. \exists ch. \text{Pt}(ch) = \text{channel} \wedge \text{S}(ch) = \text{true} \wedge \mathbf{A}(ch) = z$
8. $\forall ch. \text{Pt}(ch) = \text{channel} \wedge \text{S}(ch) = \text{true} \Rightarrow \mathbf{A}(ch) \in s$

where

- $\text{size}_h(ch)$ returns the number of the messages in the channel ch , i.e. $|h(\mathbf{A}(ch))|$
- $\text{queue}_h(ch)$ returns the messages in the channel ch , i.e. $h(\mathbf{A}(ch))$
- $\text{levels}(O)$ returns the levels of the obligations in O , i.e. $\{r \mid (a, r) \in O\}$
- $\text{Pt} = \bigsqcup_{(id, c, \xi, p, O, I, C, T) \in A} p$ and $\text{Wt} = \bigsqcup_{(id, \text{wait}(l), \xi, p, O, I, C, T) \in A} \{l\}$
- $\text{Ct} = \bigsqcup_{(id, c, \xi, p, O, I, C, T) \in A} C$ and $\text{Ot} = \bigsqcup_{(id, c, \xi, p, O, I, C, T) \in A} O$
- $\text{Tt} = \bigsqcup_{(id, c, \xi, p, O, I, C, T) \in A} T$ and $\text{lt} = \bigsqcup_{(id, c, \xi, p, O, I, C, T) \in A} I$

► **Theorem 13** (Small Steps Preserve Validity of Configurations). *Each step of the execution preserves validity of configurations.*

$$(t, h, s) \rightsquigarrow (t', h', s') \wedge \text{valid}_{n+1}(t, h, s) \Rightarrow \text{valid}_n(t', h', s')$$

Proof. By case analysis of the small step relation \rightsquigarrow .

Case $(t[id := \text{new_channel}(b); \xi], h[z := \emptyset], s) \rightsquigarrow (t[id := \text{val}(z); \xi], h[z := []], b = \text{true} ? s \cup \{z\} : s)$:
 By $\text{valid}(t[id := \text{new_channel}(b); \xi], h[z := \emptyset], s)$ we have an augmented thread set A consisting of an element $(id, \text{new_channel}(b), \xi, p, O, I, C, T)$ which satisfies all the conditions in the definition of validity of configurations, including $p, O, I, C, T \models \text{wpcx}(\text{new_channel}(b), \xi)$. $\text{valid}(t[id := \text{val}(z); \xi], h[ch := []], b = \text{true} ? s \cup \{z\} : s)$ holds because by Lemma 4 there exists an augmented thread set $A' = A - (id, \text{new_channel}(b), \xi, p, O, I, C, T) \cup (id, \text{val}(z), \xi, p[l := \text{channel}]$

, O, I, C, T) which satisfies all the conditions in the definition of validity of configurations.

Case $(t[id:=\text{send}(e_1, e_2); \xi], h[[ch]:=M], s) \rightsquigarrow (t[id:=\text{tt}; \xi], h[[e_1]:=M.[[e_2]]], s)$:

By $\text{valid}(t[id:=\text{send}(e_1, e_2); \xi], h[[ch]:=M], s)$ we have an augmented thread set A consisting of an element $(id, \text{send}(e_1, e_2), \xi, p, O, I, C, T)$ which satisfies all the conditions in the definition of validity of configurations, including $p, O, I, C, T \models \text{wpcx}(\text{send}(e_1, e_2), \xi)$. $\text{valid}(t[id:=\text{tt}; \xi], h[[e_1]:=M.[[e_2]]], s)$ holds because by Lemma 5 there exists an augmented thread set $A' = A - (id, \text{send}(e_1, e_2), \xi, p, O', I', C, T) \cup (id, \text{tt}, \xi, p_2, O \uplus M'(ch)(m), I - \{ch\}, C_2, T_2)$ which satisfies all the conditions in the definition of validity of configurations.

Case $(t[id:=\text{receive}(e); \xi], h[[e]=[m].M], s) \rightsquigarrow (t[id:=\text{val}(m); \xi], h[[e]=M], s)$:

By $\text{valid}(t[id:=\text{receive}(e); \xi], h[[e]=[m].M], s)$ we have an augmented thread set A consisting of an element $(id, \text{receive}(e), \xi, p, O, I, C, T)$ which satisfies all the conditions in the definition of validity of configurations, including $p, O, I, C, T \models \text{wpcx}(\text{receive}(e), \xi)$. $\text{valid}(t[id:=\text{val}(m); \xi], h[[e]=M], s)$ holds because by Lemma 6 there exists an augmented thread set $A' = A - (id, \text{receive}(e), \xi, p, O, I, C, T) \cup (id, \text{val}(m), \xi, p_2, O \uplus M'(ch)(m), I - \{ch\}, C_2, T_2)$ which satisfies all the conditions in the definition of validity of configurations.

Case $(t[id:=\text{fork}(c); \xi, id':=\emptyset], h, s) \rightsquigarrow (t[id:=\text{val}(\text{tt}); \xi, id':=c; \text{done}], h, s)$:

By $\text{valid}_n(t[id:=\text{fork}(c); \xi, id':=\emptyset]; \xi, h, s)$ we have an augmented thread set A consisting of an element $(id, \text{fork}(c), \xi, p, O, I, C, T)$ which satisfies all the conditions in the definition of validity of configurations, including $p, O, I, C, T \models \text{wpcx}_n(\text{fork}(c), \xi)$. $\text{valid}(t[id:=\text{val}(\text{tt}); \xi, id':=c; \text{done}], h, s)$ holds because by Lemma 7 there exists an augmented thread set $A' = A - (id, \text{fork}(c), \xi, p, O, I, C, T) \cup (id, \text{tt}, \xi, p_1, O_1, I_1, C_1, T_1) \cup (id', c, \text{done}, p_2, O_2, I_2, C_2, T_2)$ which satisfies all the conditions in the definition of validity of configurations.

Case $(t[id:=\text{let}(x, c_1, c_2); \xi], h, s) \rightsquigarrow (t[id:=c_1; \text{let}'(x, c_2, \xi)], h, s)$:

By $\text{valid}_n(t[id:=\text{let}(x, c_1, c_2); \xi], h)$ we have an augmented thread set A consisting of an element $(id, \text{let}(x, c_1, c_2), \xi, p, O, g)$ which satisfies all the conditions in the definition of validity of configurations, including $p, O, I, C, T \models \text{wpcx}_n(\text{let}(x, c_1, c_2), \xi)$. Since $\text{wpcx}_n(\text{let}(x, c_1, c_2), \xi) = \text{wp}_{n-1}(c_1, \lambda z. \text{wp}_{n-1}(c_2[z/x], Q))$, we have $p, O, I, C, T \models \text{wp}_{n-1}(c_1, \lambda z. \text{wp}_{n-1}(c_2[z/x], Q))$. Consequently $\text{valid}(t[id:=c_1; \text{let}'(x, c_2, \xi)], h, s)$ holds because there exists an augmented thread set $A' = A - (id, \text{let}(x, c_1, c_2), \xi, p, O, I, C, T) \cup (id, c_1, \text{let}'(x, c_2, \xi), p, O, I, C, T)$ which satisfies all the conditions in the definition of validity of configurations. The rest of the cases can be proved similarly. \blacktriangleleft

► **Theorem 14** (A Valid Configuration Is Not Deadlocked). *If a valid configuration has some threads then either all threads in this configuration are waiting for some server channels, or there exists a thread in this configuration which is not waiting for an empty channel.*

$$\text{valid}_n(t, h, s) \wedge \exists id. t(id) \neq \emptyset \Rightarrow \text{NotDeadlock}(t, h, s)$$

where $\text{NotDeadlock}(t, h, s) \Leftrightarrow \text{AllWaitingToServe}(t, s) \vee \exists id'. \neg \text{is_waiting}(\text{fst}(t(id')), h)$, where

■ $\text{AllWaitingToServe}(t, s) \Leftrightarrow \forall id. t(id) \neq \emptyset \Rightarrow \exists e. \text{fst}(t(id)) = \text{wait}(e) \wedge [e] \in s$

■ $\text{is_waiting}(c, h) \Leftrightarrow \exists e. c = \text{wait}(e) \wedge h([e]) = []$.

Proof. By contradiction; we assume that all threads in t are waiting for some empty channels where some of these channels are not server channels, i.e. $\forall id. \exists e. \text{fst}(t(id)) = \text{wait}(e) \wedge h(e) = [] \wedge \exists id', e'. \text{fst}(t(id')) = \text{wait}(e') \wedge [e'] \notin s$. Since (t, h, s) is a valid configuration and all threads in this configuration are waiting for a channel, there exists a set of valid augmented

22:42 Transferring Obligations Through Synchronizations

threads A from which we produce a valid bag $G = \text{valid_bag}(A)$, where valid_bag maps any element $(id, \text{wait}(e), \xi, p, O, I, C, T) \in A$ to an element $(\llbracket e \rrbracket, \text{Addresses}(O), \text{Addresses}(I))$ where $\text{Addresses}(O) = \{\mathbf{A}(o) \mid o \in O\}$. By Lemma 15, we have $G = \{\}$, implying $A = \{\}$, implying $t = \mathbf{0}$ which contradicts the hypothesis of the theorem.

Note that in the definition of validity of a configuration we also keep track of all locations whose addresses are allocated, which makes it possible to provide the functions R, M^r , and S , mapping channel addresses to their ghost information, for Lemma 15. Additionally, the hypotheses H_2, H_3, H_4 , and $H_5 \vee H_6$ in Lemma 15 are met as follows. For each element $(id, \text{wait}(e), \xi, p, O, I, C, T) \in A$ we have $p, O, I, C, T \models \text{wpcx}(\text{wait}(e), \xi)$, which implies $0 < C(\llbracket e \rrbracket)$ and there exists a channel ch with address $\llbracket e \rrbracket$ such that $ch \prec O$ (which implies H_2) and $ch \prec^r I$ (which implies H_3), and $S(ch) = \text{true} \Rightarrow O = \{\} \wedge \forall o_1. 0 < I(o_1) \Rightarrow o_1 = ch$ (which implies H_4). Additionally, by $0 < C(\llbracket e \rrbracket)$ we have $0 < \text{Ct}(\llbracket e \rrbracket)$, which (by 4.c in validity of configurations) implies if $S(ch) = \text{false}$ either 1) $0 < \text{Ot}(ch)$ (which implies H_5), or 2) there exists a message m in a channel ch' through which an obligation of ch is transferred, i.e. $0 < \text{size}(ch')$ and $m \in \text{queue}_h(ch')$ and $0 < M'(ch')(m)(ch)$. By $m \in \text{queue}_h(ch')$ and 4.b we have $\text{levels}(M'(ch')(m)) \subseteq M'(ch')$, which by $\{ch\} \in M'(ch')(m)$ implies $R(ch) \in M'(ch')$. Additionally, by $M'(ch') \neq \{\}$, and $0 < \text{size}(ch')$, and 4.a we have $0 < \text{lt}(ch')$ (which implies H_6). \blacktriangleleft

Lemma 15 ensures that in any state of the execution if all the desired invariants are respected then it is impossible that all threads of the program are waiting for some empty channels where some of these channels are not server channels. In this lemma G is a bag of waitable object-obligations-importers triples such that each element t of G is associated with a thread in a state of the execution, where the first element of t is associated with the address of the object for which t is waiting, the second element is associated with the addresses of obligations of t , and the third element is associated with the addresses of importers of t .

► **Lemma 15** (A Valid Bag of Augmented Threads Is Not Deadlocked).

$$\begin{aligned} \forall G : & \text{Bags}(\text{Addresses} \times \text{Bags}(\text{Addresses}) \times \text{Bags}(\text{Obligations})), \\ & R : \text{Addresses} \rightarrow \text{Levels}, \\ & M^r : \text{Addresses} \rightarrow \text{Bags}(\text{Levels}), \\ & S : \text{Addresses} \rightarrow \text{Booleans}. \\ & H_1 \wedge \forall (o, O, I) \in G. H_2 \wedge H_3 \wedge H_4 \wedge (H_5 \vee H_6) \Rightarrow G = \{\} \end{aligned}$$

where

- $H_1 : \exists (o, O, I) \in G. S(o) = \text{false}$
- $H_2 : o \prec O$
- $H_3 : o \prec^r I$
- $H_4 : S(o) = \text{true} \Rightarrow O = \{\} \wedge \forall o_1. 0 < I(o_1) \Rightarrow o_1 = o$
- $H_5 : S(o) = \text{false} \Rightarrow 0 < \text{Ot}(o)$
- $H_6 : S(o) = \text{false} \Rightarrow \exists o_1. R(o) \in M^r(o_1) \wedge \text{Wt}(o_1) = 0 \wedge 0 < \text{lt}(o_1)$
 where $\text{Wt} = \bigsqcup_{(o, O, I) \in G} \{\mathbf{o}\}$ and $\text{Ot} = \bigsqcup_{(o, O, I) \in G} O$ and $\text{lt} = \bigsqcup_{(o, O, I) \in G} I$

Proof. By H_1 we know $\exists (o_m, O_1, I_1) \in G$ where $S(o_m) = \text{false}$ and $\forall (o, O, I) \in G. S(o) = \text{false} \Rightarrow R(o_m) \leq R(o)$. By $(H_5 \vee H_6)$ there are two cases: 1) $\exists (o_2, \{\mathbf{o}_m\} \sqcup O_2, I_2) \in G$, or 2) $\exists (o_3, O_3, \{\mathbf{o}_1\} \sqcup I_3) \in G. R(o_m) \in M^r(o_1) \wedge \text{Wt}(o_1) = 0$. In the first case by H_4 we have $S(o_2) = \text{false}$, which implies $R(o_m) \leq R(o_2)$, that contradicts the hypothesis H_2 , i.e. $o_2 \prec \{\mathbf{o}_m\} \sqcup O_2$. In the second case by $\text{Wt}(o_1) = 0$ we have $o_1 \neq o_3$ (because $0 < \text{Wt}(o_3)$), and consequently by H_4 we have $S(o_3) = \text{false}$, which implies $R(o_m) \leq R(o_3)$, that contradicts the hypothesis H_3 , i.e. $o_3 \prec^r \{\mathbf{o}_1\} \sqcup I_3$ (because $R(o_m) \in M^r(o_1)$). \blacktriangleleft

► **Theorem 16** (The Initial Configuration Is Valid).

$$\text{correct}(\text{obs}(\{\}, \{\}), c, \lambda_.\text{obs}(\{\}, \{\})) \Rightarrow \forall n, id. \text{valid}_n(\mathbf{0}[id:=c; \text{done}], \mathbf{0}, [])$$

Proof. The goal is achieved because there exists an augmented thread set $A = [(id, c, \text{done}, \mathbf{0}, \theta, \theta, \theta)]$, such that all the conditions in the definition of validity of configurations are met, where $\mathbf{0} = \lambda_.\emptyset$ and $\theta = \lambda_.\mathbf{0}$. ◀

C.6 An Example Proof

In this section we show how the program in Figure 3 can be verified using the proof rules in Figure 31, as shown in Figure 32.

D Transferring Obligations Through Notifications: Soundness Proof

In this appendix we provide a formalization and soundness proof, machine-checked with Coq¹¹, for the approach introduced in Section 3. However, unfortunately, there are a few technical differences between this formalization and the system of Section 3 such that in this formalization the ghost information, such as level and transferred permissions and obligations, are associated with lock and condition variable addresses via the `lock` and `cond` permissions rather than via global functions¹². The proof rules associated with this formalization and the verification of the program in Figure 16, proved using these rules, are shown in Sections D.4 and D.6, respectively.

D.1 Syntax and Semantics of Programs

We define the syntax of our programming language as indicated in Figure 33. In this syntax an arithmetic expression, e , can be an integer value, z , a variable, x , an addition of two expressions, or a negation of an expression. An integer value can be substituted for a free variable in an expression or a command, and each expression can be evaluated to an integer value, as shown in Figure 34. Commands include commands `val(e)` which simply yield the value of e as their result and have no side effects, memory allocations¹³, memory reads, memory writes, conditionals, loops, parallel composition, sequential composition, lock creations, lock acquisitions, lock releases, condition variable creations, waits, and notifications. We also define some extra commands `waiting4lock`, indicating that the related thread is waiting for a lock, and `waiting4cvar`, indicating that the related thread has executed `wait`; these are not supposed to appear in the source program and appear only during execution. Additionally, instead of defining all the ghost commands introduced in Section 3, we define a single ghost command `nop` which is inserted into the program for verification purposes and has no effect on the program's behavior. The small step semantics, defined in Figure 36, relates two *configurations*, defined in Figure 35.

22:44 Transferring Obligations Through Synchronizations

```

ob(a) ::= (a, 1, {1}, false)
ob'(a) ::= (a, 1, {}, false)
loc(a) ::= (ob(a), (Mch, []), λm. {ob'(snd(m))})
loc'(a) ::= (ob'(a), (Mch', []), λ_. {})
pt(Mch, args) ::= λm. channel(loc'(snd(m)))
pt(Mch', args) ::= λm. true

routine server(channel a){
req : {obs({}, {ob(a)}) * channel(loc(a)) * credit(a)}
  (req, a') := receive(ch);
  {obs({ob'(a')}, {} * channel(loc'(a'))}
  result := process(req);
  send(a', result)
ens : {obs({}, {})}

routine client(channel a){
req : {obs({ob(a)}, {}) * channel(loc(a)) * trandit(a)}
  a' := new_channel();
  {obs({ob(a)}, {}) * channel(loc'(a')) * trandit(a)}
  nop; //Rule CREDIT
  {obs({ob(a), ob'(a')}, {}) * trandit(a) * credit(a')}
  send(a, (request(), a'));
  {obs({}, {}) * credit(a')}
  receive(a')
ens : {obs({}, {})}

routine main(){
req : {obs({}, {})}
  a := new_channel();
  {obs({}, {}) * channel(loc(a))}
  nop; //Rule CREDIT
  {obs({ob(a)}, {}) * credit(a)}
  nop; //Rule TRANDIT
  {obs({ob(a)}, {ob(a)}) * credit(a) * trandit(a)}
  fork(
    {obs({}, {ob(a)}) * credit(a)}
    server(a)
    {obs({}, {})});
  {obs({ob(a)}, {}) * trandit(a)}
  client(a)
ens : {obs({}, {})}

```

■ **Figure 32** Verification of the program in Figure 3 using the rules in Figure 31.

$c \in \text{Commands}, e \in \text{Expressions}, z \in \mathbb{Z}, x \in \text{Variables}$
 $e ::= z \mid x \mid e_1 + e_2 \mid -e$
 $c ::= \text{val}(e) \mid \text{new_int}(z) \mid \text{lookup}(e) \mid \text{mutate}(e_1, e_2)$
 $\quad \mid \text{if}(c, c_1, c_2) \mid \text{while}(c, c_1) \mid \text{let}(x, c_1, c_2) \mid \text{fork}(c) \mid \text{new_lock} \mid \text{acquire}(e) \mid \text{release}(e)$
 $\quad \mid \text{new_cvar} \mid \text{wait}(e_1, e_2) \mid \text{notify}(e) \mid \text{notifyAll}(e)$
 $\quad \mid \text{waiting4lock}(e) \mid \text{waiting4cvar}(e_1, e_2) \mid \text{nop}$

■ **Figure 33** Syntax of the programming language

$\llbracket \cdot \rrbracket \in \text{Expressions} \rightarrow \mathbb{Z}$

$\llbracket z \rrbracket = z$

$\llbracket x \rrbracket = 0$

$\llbracket e_1 + e_2 \rrbracket = \llbracket e_1 \rrbracket + \llbracket e_2 \rrbracket$

$\llbracket -e \rrbracket = -\llbracket e \rrbracket$

$z[z'/x] = z$

$x[z/x'] = x \quad \text{if } x \neq x'$

$x[z/x] = z$

$(e_1 + e_2)[z/x] = e_1[z/x] + e_2[z/x]$

$\text{val}(e)[z/x] = \text{val}(e[z/x])$

$\text{let}(x, c_1, c_2)[z/x] = \text{let}(x, c_1, c_2)$

$\text{let}(x, c_1, c_2)[z/x'] = \text{let}(x, c_1[z/x'], c_2[z/x']) \quad \text{if } x \neq x'$

$\text{fork}(c)[z/x] = \text{fork}(c[z/x])$

$\text{new_lock}[z/x] = \text{new_lock}$

$\text{new_int}(z')[z/x] = \text{new_int}(z')$

$\text{lookup}(e)[z/x] = \text{lookup}(e[z/x])$

$\text{mutate}(e_1, e_2)[z/x] = \text{mutate}(e_1[z/x], e_2[z/x])$

$\text{acquire}(e)[z/x] = \text{acquire}(e[z/x])$

$\text{release}(e)[z/x] = \text{release}(e[z/x])$

$\text{new_cvar}[z/x] = \text{new_cvar}$

$\text{wait}(e_1, e_2)[z/x] = \text{wait}(e_1[z/x], e_2[z/x])$

$\text{notify}(e)[z/x] = \text{notify}(e[z/x])$

$\text{notifyAll}(e)[z/x] = \text{notifyAll}(e[z/x])$

$\text{waiting4lock}(e)[z/x] = \text{waiting4lock}(e[z/x])$

$\text{waiting4cvar}(e_1, e_2)[z/x] = \text{waiting4cvar}(e_1[z/x], e_2[z/x])$

$\text{while}(c, c_1)[z/x] = \text{while}(c[z/x], c_1[z/x])$

$\text{if}(c, c_1, c_2)[z/x] = \text{if}(c[z/x], c_1[z/x], c_2[z/x])$

$\text{nop}[z/x] = \text{nop}$

■ **Figure 34** Evaluation of expressions and substitution of expressions and commands

22:46 Transferring Obligations Through Synchronizations

$$\begin{aligned}
\text{Addresses} &= \mathbb{Z} \\
\text{ThreadId}s &= \mathbb{Z} \\
h \in \text{Heaps} &= \text{Addresses} \rightarrow \mathbb{Z} \\
\xi \in \text{Contexts} &::= \text{done} \mid \text{let}'(x, c, \xi) \mid \text{if}'(c_1, c_2, \xi) \\
\theta \in \text{ThreadConfigurations} &::= (c; \xi) \\
t \in \text{ThreadTables} &= \text{ThreadId}s \rightarrow \text{ThreadConfigurations} \\
\kappa \in \text{Configurations} &= \text{Heaps} \times \text{ThreadTables}
\end{aligned}$$

■ **Figure 35** Configurations

$$\begin{aligned}
&(t[id:=\text{new_int}(n); \xi], h[z\dots z+n-1:=\emptyset]) \rightsquigarrow (t[id:=\text{val}(z); \xi], h[z\dots z+n-1:=0]) \\
&(t[id:=\text{lookup}(e); \xi], h[\llbracket e \rrbracket:=z]) \rightsquigarrow (t[id:=\text{val}(z); \xi], h[\llbracket e \rrbracket:=z]) \\
&(t[id:=\text{mutate}(e_1, e_2); \xi], h) \rightsquigarrow (t[id:=\text{tt}; \xi], h[\llbracket e_1 \rrbracket:=\llbracket e_2 \rrbracket]) \\
&(t[id:=\text{if}(c, c_1, c_2); \xi], h) \rightsquigarrow (t[id:=c; \text{if}'(c_1, c_2, \xi)]; h) \\
&(t[id:=\text{val}(e); \text{if}'(c_1, c_2, \xi)]; h) \rightsquigarrow (t[id:=c_1; \xi]; h) \quad \text{if } 0 < \llbracket e \rrbracket \\
&(t[id:=\text{val}(e); \text{if}'(c_1, c_2, \xi)]; h) \rightsquigarrow (t[id:=c_2; \xi]; h) \quad \text{if } \llbracket e \rrbracket \leq 0 \\
&(t[id:=\text{while}(c, c_1); \xi], h) \rightsquigarrow (t[id:=\text{if}(c, \text{let}(x, c_1, \text{while}(c, c_1)), \text{tt}); \xi]; h) \\
&\quad \text{where } x \text{ is not free in } c \text{ and } c_1 \\
&(t[id:=\text{fork}(c); \xi, id':=\emptyset], h) \rightsquigarrow (t[id:=\text{tt}; \xi, id':=c; \text{done}]; h) \\
&(t[id:=\text{let}(x, c_1, c_2); \xi], h) \rightsquigarrow (t[id:=c_1; \text{let}'(x, c_2, \xi)]; h) \\
&(t[id:=\text{val}(e); \text{let}'(x, c, \xi)]; h) \rightsquigarrow (t[id:=c[\llbracket e \rrbracket/x]; \xi]; h) \\
&(t[id:=\text{val}(e); \text{done}]; h) \rightsquigarrow (t[id:=\emptyset]; h) \\
&(t[id:=\text{new_lock}; \xi], h[z:=\emptyset]) \rightsquigarrow (t[id:=\text{val}(z); \xi], h[z:=1]) \\
&(t[id:=\text{acquire}(e); \xi], h[\llbracket e \rrbracket:=1]) \rightsquigarrow (t[id:=\text{tt}; \xi], h[\llbracket e \rrbracket:=0]) \\
&(t[id:=\text{acquire}(e); \xi], h[\llbracket e \rrbracket:=0]) \rightsquigarrow (t[id:=\text{waiting4lock}(e); \xi], h[\llbracket e \rrbracket:=0]) \\
&(t[id:=\text{waiting4lock}(e); \xi], h[\llbracket e \rrbracket:=1]) \rightsquigarrow (t[id:=\text{tt}; \xi], h[\llbracket e \rrbracket:=0]) \\
&(t[id:=\text{release}(e); \xi], h) \rightsquigarrow (t[id:=\text{tt}; \xi], h[\llbracket e \rrbracket:=1]) \\
&(t[id:=\text{new_cvar}; \xi], h[z:=\emptyset]) \rightsquigarrow (t[id:=\text{val}(z); \xi], h[z:=0]) \\
&(t[id:=\text{wait}(e_1, e_2); \xi], h) \rightsquigarrow (t[id:=\text{waiting4cvar}(e_1, e_2); \xi], h[\llbracket e_2 \rrbracket:=1]) \\
&(t[id:=\text{notify}(e); \xi, id':=\text{waiting4cvar}(e_1, e_2); \xi'], h) \rightsquigarrow \\
&\quad (t[id:=\text{val}(\text{tt}); \xi, id':=\text{waiting4lock}(e_2); \xi'], h) \quad \text{if } \llbracket e \rrbracket = \llbracket e_1 \rrbracket \\
&(t[id:=\text{notify}(e); \xi], h) \rightsquigarrow (t[id:=\text{tt}; \xi], h) \quad \text{if } \text{nowaiting}(\llbracket e \rrbracket, t) \\
&(t[id:=\text{notifyAll}(e); \xi], h) \rightsquigarrow (\text{wkup}(\llbracket e \rrbracket, t[id:=\text{val}(\text{tt}); \xi]), h) \\
&(t[id:=\text{nop}; \xi], h) \rightsquigarrow (t[id:=\text{tt}; \xi], h)
\end{aligned}$$

where

$$\text{wkup}(z, t) = \lambda id. \begin{cases} \text{waiting4lock}(l); \xi & \text{if } t(id) = \text{waiting4cvar}(v, l); \xi \wedge \llbracket v \rrbracket = z \\ t(id) & \text{otherwise} \end{cases}$$

$\text{nowaiting}(z, t) \Leftrightarrow \exists id, \xi, l, v. \llbracket v \rrbracket = z \wedge t(id) = \text{waiting4cvar}(v, l); \xi$

■ **Figure 36** Semantics of programs, where tt stands for val(0).

$$\begin{aligned}
& \text{Bags}(A) = A \rightarrow \mathbb{N} \\
& Wt, Ot \in \text{Bags}(\mathbb{Z}) \\
& \text{Indexes} = \mathbb{Z} \\
& \text{Arguments} = \mathbb{Z} \\
& r \in \text{Levels} = \mathbb{R} \\
& o \in \text{Obligations} = \text{Addresses} \times \text{Levels} \times \text{Addresses} \\
& l \in \text{Locations} = \text{Obligations} \times (\text{Indexes} \times \text{Lists}(\text{Arguments})) \times \\
& \quad (\text{Indexes} \times \text{Lists}(\text{Arguments})) \times \text{Bags}(\text{Obligations}) \\
& O \in \text{Bags}(\text{Obligations}) \\
& b \in \text{Booleans} \\
& \hat{v} \in \text{AValues} ::= z \mid r \mid b \mid l \mid o \mid O \\
& \alpha \in \text{AValues} \rightarrow \text{Assertions} \\
& \pi \in \text{Fractions} \\
& a \in \text{Assertions} ::= l \xrightarrow{\pi} z \mid \text{unlock}(l, Wt, Ot) \mid \text{lock}(l) \mid \text{locked}(l, Wt, Ot) \mid \text{ucond}(l) \mid \text{cond}(l) \\
& \quad \mid \text{obs}(O) \mid \text{ctr}(z, n) \mid \text{tic}(z) \\
& \quad \mid b \mid a_1 \wedge a_2 \mid a_1 \vee a_2 \mid a_1 * a_2 \mid a_1 \multimap a_2 \mid \forall \alpha \mid \exists \alpha \\
\\
& \text{pt} : \text{PredicateTables} = \text{Indexes} \rightarrow \text{Lists}(\text{Arguments}) \rightarrow \\
& \quad \text{Bags}(\text{Obligations}) \rightarrow \text{Bags}(\text{Obligations}) \rightarrow \text{Assertions} \\
& \text{O} : \text{Locations} \rightarrow \text{Obligations} , \text{ where } \text{O}((A, R, L), I, M, M') = (A, R, L) \\
& \text{A} : \text{Locations} \rightarrow \text{Addresses} , \text{ where } \text{A}((A, R, L), I, M, M') = A \\
& \text{R} : \text{Locations} \rightarrow \text{Levels} , \text{ where } \text{R}((A, R, L), I, M, M') = R \\
& \text{L} : \text{Locations} \rightarrow \text{Addresses} , \text{ where } \text{L}((A, R, L), I, M, M') = L \\
& \text{I} : \text{Locations} \rightarrow \text{Bags}(\text{Obligations}) \rightarrow \text{Bags}(\text{Obligations}) \rightarrow \text{Assertions} \\
& \quad \text{where } \text{I}((A, R, L), I, M, M') = \text{pt}(\text{fst}(I), \text{snd}(I)) \\
& \text{M} : \text{Locations} \rightarrow \text{Assertions}, \\
& \quad \text{where } \text{M}((A, R, L), I, M, M') = \text{pt}(\text{fst}(M), \text{snd}(M), \{\}, \{\}) \\
& \text{M}' : \text{Locations} \rightarrow \text{Bags}(\text{Obligations}) , \text{ where } \text{M}'((A, R, L), I, M, M') = M'
\end{aligned}$$

■ **Figure 37** Syntax of assertions

22:48 Transferring Obligations Through Synchronizations

$$\begin{aligned}
k &\in \text{Knowledge} ::= \text{cell}(\pi, z) \mid \text{unlock}(Wt, Ot) \mid \text{lock} \mid \text{locked}(Wt, Ot) \mid \text{ucond} \mid \text{cond} \\
p &\in \text{PermissionHeaps} = \text{Locations} \rightarrow \text{Knowledge} \\
\text{GhostIdentifications} &= \mathbb{Z} \\
gv &\in \text{GhostValues} ::= \text{Option}(\mathbb{N}) \times \mathbb{N} \\
g &\in \text{GhostHeaps} = \text{GhostIdentifications} \rightarrow \text{GhostValues} \\
\text{Option}(A) &::= s \mid \emptyset, \text{ where } s \in A \\
\tilde{O} &\in \text{Option}(\text{Bags}(\text{Obligations})) \\
\\
p, \tilde{O}, g &\models l \xrightarrow{\pi} z \Leftrightarrow p(l) = \text{cell}(\pi, z) \\
p, \tilde{O}, g &\models \text{unlock}(l, Wt, Ot) \Leftrightarrow p(l) = \text{unlock}(Wt, Ot) \\
p, \tilde{O}, g &\models \text{lock}(l) \Leftrightarrow p(l) = \text{lock}() \\
p, \tilde{O}, g &\models \text{locked}(l, Wt, Ot) \Leftrightarrow p(l) = \text{locked}(Wt, Ot) \\
p, \tilde{O}, g &\models \text{ucond}(l) \Leftrightarrow p(l) = \text{ucond} \\
p, \tilde{O}, g &\models \text{cond}(l) \Leftrightarrow p(l) = \text{cond} \\
p, \tilde{O}, g &\models \text{obs}(O) \Leftrightarrow \tilde{O} = O \\
p, \tilde{O}, g &\models \text{ctr}(z, n) \Leftrightarrow \exists n_1. g(z) = (n, n_1) \\
p, \tilde{O}, g &\models \text{tic}(z) \Leftrightarrow \exists n, \tilde{n}. g(z) = (\tilde{n}, n+1) \\
p, \tilde{O}, g &\models b \Leftrightarrow b = \text{true} \\
p, \tilde{O}, g &\models a_1 \wedge a_2 \Leftrightarrow p, \tilde{O}, g \models a_1 \wedge p, \tilde{O}, g \models a_2 \\
p, \tilde{O}, g &\models a_1 \vee a_2 \Leftrightarrow p, \tilde{O}, g \models a_1 \vee p, \tilde{O}, g \models a_2 \\
p, \tilde{O}, g &\models a_1 * a_2 \Leftrightarrow \exists p_1, p_2, \tilde{O}_1, \tilde{O}_2, g_1, g_2. p = p_1 \uplus p_2 \wedge \tilde{O} = \tilde{O}_1 \uplus \tilde{O}_2 \wedge g = g_1 \uplus g_2 \wedge \\
&\quad p_1, \tilde{O}_1, g_1 \models a_1 \wedge p_2, \tilde{O}_2, g_2 \models a_2 \\
p, \tilde{O}, g &\models a_1 * a_2 \Leftrightarrow \forall p_1, \tilde{O}_1, g_1. p_1, \tilde{O}_1, g_1 \models a_1 \Rightarrow \\
&\quad \forall p_2, \tilde{O}_2, g_2. p_2 = p \uplus p_1 \wedge g_2 = g \uplus g_1 \wedge \tilde{O}_2 = \tilde{O} \uplus \tilde{O}_1 \Rightarrow p_2, \tilde{O}_2, g_2 \models a_2 \\
p, \tilde{O}, g &\models \forall \alpha \Leftrightarrow \forall \hat{v} \in A\text{Values}. p, \tilde{O}, g \models \alpha(\hat{v}) \\
p, \tilde{O}, g &\models \exists \alpha \Leftrightarrow \exists \hat{v} \in A\text{Values}. p, \tilde{O}, g \models \alpha(\hat{v}) \\
\\
a_1 \vdash a_2 &\Leftrightarrow (\forall p, \tilde{O}, g. p, \tilde{O}, g \models a_1 \Rightarrow p, \tilde{O}, g \models a_2)
\end{aligned}$$

■ **Figure 38** Satisfaction relation

$$\begin{aligned}
& O_1, O_2 \in \text{Bags}(A) \\
& \tilde{O}_1, \tilde{O}_2 \in \text{Option}(A) \\
& gv_1, gv_2 \in \text{GhostValues} \\
& g_1, g_2 \in \text{GhostIdentifications} \rightarrow \text{GhostValues} \\
& k_1, k_2 \in \text{Knowledge} \\
& p_1, p_2 \in \text{PermissionHeaps} \\
\\
& O_1 \uplus O_2 = \lambda v. O_1(v) + O_2(v) \\
\\
& \tilde{O}_1 \uplus \tilde{O}_2 = \begin{cases} \tilde{O}_1 & \text{if } \tilde{O}_2 = \emptyset \\ \tilde{O}_2 & \text{if } \tilde{O}_1 = \emptyset \\ \text{undefined} & \text{otherwise} \end{cases} \\
\\
& \tilde{O}_1 \uplus \tilde{O}_2 = \begin{cases} \tilde{O}_1 & \text{if } \tilde{O}_2 = \emptyset \\ \tilde{O}_2 & \text{if } \tilde{O}_1 = \emptyset \\ \tilde{O}_1 \uplus \tilde{O}_2 & \text{otherwise} \end{cases} \\
\\
& gv_1 \uplus gv_2 = \begin{cases} (\tilde{m}_1 \uplus \tilde{m}_2, n_1 + n_2) & \text{if } gv_1 = (\tilde{m}_1, n_1) \wedge gv_2 = (\tilde{m}_2, n_2) \wedge \\ & (\tilde{m}_1 \uplus \tilde{m}_2 = n \Rightarrow n \geq n_1 + n_2) \\ \text{undefined} & \text{otherwise} \end{cases} \\
\\
& g_1 \uplus g_2 = \begin{cases} \lambda l. g_1(l) \uplus g_2(l) & \text{if } \exists g. \forall l. g(l) = g_1(l) \uplus g_2(l) \\ \text{undefined} & \text{otherwise} \end{cases} \\
\\
& k_1 \uplus k_2 = \begin{cases} \text{cell}(\pi + \pi', z) & \text{if } k_1 = \text{cell}(\pi, z) \wedge k_2 = \text{cell}(\pi', z) \wedge \pi + \pi' \leq 1 \\ \text{lock} & \text{if } k_1 = \text{lock} \wedge k_2 = \text{lock} \\ \text{locked}(Wt, Ot) & \text{if } k_1 = \text{lock} \wedge k_2 = \text{locked}(Wt, Ot) \\ \text{locked}(Wt, Ot) & \text{if } k_1 = \text{locked}(Wt, Ot) \wedge k_2 = \text{lock} \\ \text{cond} & \text{if } k_1 = \text{cond} \wedge k_2 = \text{cond} \\ \text{undefined} & \text{otherwise} \end{cases} \\
\\
& p_1 \uplus p_2 = \begin{cases} \lambda l. p_1(l) \uplus p_2(l) & \text{if } \exists p. \forall l. p(l) = p_1(l) \uplus p_2(l) \\ \text{undefined} & \text{otherwise} \end{cases}
\end{aligned}$$

■ **Figure 39** Operations on ghost resources

$$\begin{aligned}
& \text{pheap_heap}(p, h) \Leftrightarrow \\
& (\forall z. (\forall l. A(l) = z \Rightarrow p(l) = \emptyset) \Rightarrow h(z) = \emptyset) \text{ and} \\
& \forall l. \\
& \quad p(l) = \emptyset \text{ or} \\
& \quad \forall z. (\exists \pi. p(l) = \text{cell}(\pi, z) \Rightarrow h(A(l)) = z) \text{ and} \\
& \quad (\exists O_1, O_2. p(l) = \text{unlock}(O_1, O_2) \Rightarrow h(A(l)) = 1) \text{ and} \\
& \quad p(l) = \text{lock} \Rightarrow h(A(l)) = 1 \text{ and} \\
& \quad (\exists O_1, O_2. p(l) = \text{locked}(O_1, O_2) \Rightarrow h(A(l)) = 0) \text{ and} \\
& \quad p(l) = \text{cond} \Rightarrow h(A(l)) \neq \emptyset \text{ and} \\
& \quad p(l) = \text{ucond} \Rightarrow h(A(l)) \neq \emptyset
\end{aligned}$$

■ **Figure 40** Permission heaps corresponding to concrete heaps

D.2 Syntax and Semantics of Assertions

The syntax of assertions is defined in Figure 37¹⁴. Note that a *location* l of an object o consists of the *obligation* of o (if o is a lock or a CV); the lock invariant of o (if o is a lock), denoted by $l(l)$; and the permissions and the obligations which are transferred through a notification on o (if o is a CV), denoted by $M(l)$ and $M'(l)$ respectively. Also note that permissions described by invariants of locks as well as permissions which are transferred through notifications are specified through an index (as well as the required arguments) pointing to a table in which each element is a function that given a list of arguments returns an assertion. This makes it possible to quantify over locations in assertions¹⁵. The obligation of a location l , denoted by $O(l)$, consists of the address of that location, denoted by $A(l)$, as well as other ghost information such as the level of l , denoted by $R(l)$; and the lock associated with l (if l is the location of a CV), denoted by $L(l)$.

These assertions describe some ghost resources, namely p , O , and g that keep track of heap locations, obligations, and ghost counters, respectively, shown in Figure 38, where some operations and relations defined on these resources are shown in Figures 39, 40.

D.3 Weakest Precondition of Commands

The weakest precondition of a command c for $n > 0$ steps w.r.t. a postcondition Q (with a given predicate table, specified by pt), denoted by $\text{wp}_{n,pt}(c, Q)$ is defined in Figures 41 and 42. Note that $\text{wp}(c, Q)_{0,pt} = \text{true}$. Also note that for the sake of simplicity the index pt is elided. Having this definition, we define the weakest precondition of a context and the

¹¹ The soundness proof of the second mechanism, machine-checked in Coq, can be found in [17].

¹² Note that one way to formalize the precise approach of Section 3 would be to define assertions as functions from ghost information to separating conjunctions of chunks. In the soundness proof, one would track these as partial functions whose domain is the set of allocated addresses. The functions passed into the assertions would be *totalizations* of these partial functions. An assertion is true if it is true for all totalizations of the functions.

¹³ Note that `new_int(n)` allocates n consecutive memory locations and returns the address of the first one.

¹⁴ Note that we use a shallow embedding: assertions have no variables; to model quantifications, we use meta-level functions from values to assertions.

¹⁵ An alternative approach is to use a step-indexed domain of assertions, as in Iris [27]. There, \blacktriangleright *Assertions* could be used instead of *Indexes* \times *Lists(Arguments)*, where \blacktriangleright is Iris's *guard* for *guarded recursive definitions*.

$$\begin{aligned}
& \text{wp} \in \text{WeakestPreconditions} = \\
& \quad \text{Commands} \rightarrow (\mathbb{Z} \rightarrow \text{Assertions}) \rightarrow \mathbb{N} \rightarrow \text{PredicateTables} \rightarrow \text{Assertions} \\
& \text{wp}_n(\text{val}(e), Q) = Q(\llbracket e \rrbracket) \\
& \text{wp}_n(\text{new_lock}, Q) = \forall z. \exists r. \text{ulock}(((z, r, z), (0, []), (0, []), \{\}, \{\}, \{\}) \multimap Q(z)) \\
& \text{wp}_n(\text{acquire}(e), Q) = \exists O, l. (\text{lock}(l) * \text{obs}(O) \wedge O(l) \prec O \wedge A(l) = \llbracket e \rrbracket) * \\
& \quad (\forall Wt, Ot. (\text{obs}(O \uplus \{O(l)\}) * \text{locked}(l, Wt, Ot) * l(l)(Wt, Ot)) \multimap Q(\text{tt})) \\
& \text{wp}_n(\text{waiting4lock}(e), Q) = \text{wp}_n(\text{acquire}(e), Q) \\
& \text{wp}_n(\text{release}(e), Q) = \exists Wt, Ot, O, l. (\text{locked}(l, Wt, Ot) * l(l)(Wt, Ot) * \text{obs}(O \uplus \{O(l)\}) \wedge \\
& \quad A(l) = \llbracket e \rrbracket) * ((\text{lock}(l) * \text{obs}(O)) \multimap Q(\text{tt})) \\
& \text{wp}_n(\text{new_cvar}, Q) = \forall z. \exists r. \text{ucond}(((z, r, z), (0, []), (0, []), \{\})) \multimap Q(z) \\
& \text{wp}_n(\text{wait}(e_1, e_2), Q) = \exists Wt, Ot, O, v, l. (\text{cond}(v) * \text{locked}(l, Wt, Ot) * \\
& \quad l(l)(Wt \uplus \{A(v)\}, Ot) * \text{obs}(O \uplus \{O(l)\}) \wedge L(v) = A(l) \wedge O(v) \prec O \wedge O(l) \prec O \uplus M'(v) \wedge \\
& \quad \text{enoughObs}(v, Wt \uplus \{A(v)\}, Ot) \wedge A(v) = \llbracket e_1 \rrbracket \wedge A(l) = \llbracket e_2 \rrbracket) * \\
& \quad (\forall Wt', Ot'. (\text{cond}(v) * \text{locked}(l, Wt', Ot') * l(l)(Wt', Ot') * \\
& \quad \text{obs}(O \uplus \{O(l)\} \uplus M'(v)) * M(v)) \multimap Q(\text{tt})) \\
& \text{wp}_n(\text{waiting4cvar}(e_1, e_2), Q) = \exists O, v, l. (\text{cond}(v) * \text{lock}(l) * \text{obs}(O) \wedge \\
& \quad O(v) \prec O \wedge O(l) \prec O \uplus M' \wedge L(v) = A(l) \wedge A(v) = \llbracket e_1 \rrbracket \wedge A(l) = \llbracket e_2 \rrbracket) * \\
& \quad (\forall Wt, Ot. (\text{cond}(v) * \text{locked}(l, Wt, Ot) * l(l)(Wt, Ot) * \\
& \quad \text{obs}(O \uplus \{O(l)\} \uplus M'(v)) * M(v)) \multimap Q(\text{tt})) \\
& \text{wp}_n(\text{notify}(e), Q) = \exists Wt, Ot, O, v. (\text{cond}(v) * \text{locked}(L(v), Wt, Ot) * \\
& \quad \text{obs}(O \uplus (O \prec Wt(A(v)) ? M'(v) : \{\})) * (Wt(A(v)) = 0 \vee M(v)) \wedge A(v) = \llbracket e \rrbracket) * \\
& \quad ((\text{cond}(v) * \text{locked}(L(v), Wt - \{A(v)\}, Ot) * \text{obs}(O)) \multimap Q(\text{tt})) \\
& \text{wp}_n(\text{notifyAll}(e), Q) = \exists Wt, Ot, O, v, l. (\text{cond}(v) * \text{locked}(l, Wt, Ot) * (\prod_{i=1}^{Wt(v)} M(v)) \wedge \\
& \quad M'(v) = \{\} \wedge A(v) = \llbracket e \rrbracket \wedge L(v) = A(l)) * ((\text{cond}(v) * \text{locked}(l, Wt[A(v) := 0], Ot)) \multimap Q(\text{tt})) \\
& \text{wp}_n(\text{let}(x, c_1, c_2), Q) = \text{wp}_{n-1}(c_1, \lambda z. \text{wp}_{n-1}(c_2[z/x], Q)) \\
& \text{wp}_n(\text{fork}(c), Q) = \exists O_1, O_2. \text{obs}(O_1 \uplus O_2) * (\text{obs}(O_1) \multimap Q(\text{tt})) * (\text{obs}(O_2) \multimap \\
& \quad \text{wp}_{n-1}(c, \lambda _ . \text{obs}(\{\}))) \\
& \text{wp}_n(\text{if}(c, c_1, c_2), Q) = \text{wp}_{n-1}(c, (\lambda z. 0 < z ? \text{wp}_{n-1}(c_1, Q) : \text{wp}_{n-1}(c_2, Q))) \\
& \text{wp}_n(\text{while}(c, c_1), Q) = \text{wp}_{n-1}(c, (\lambda z. z \leq 0 ? Q(\text{tt}) : \\
& \quad \text{wp}_{n-1}(c_1, (\lambda _ . \text{wp}_{n-1}(\text{while}(c, c_1), Q))))))
\end{aligned}$$

■ **Figure 41** Weakest precondition, where tt stands for 0 (part one of two).

22:52 Transferring Obligations Through Synchronizations

$\text{wp}_{n,pt}(\text{nop}) =$	
as g_initl	$(\exists Wt, Ot, O, a, r, I. \text{ulock}(((a, r, a), (0, []), (0, []), \{\}), Wt, Ot) * pt(I)(Wt, Ot) * \text{obs}(O) * ((\text{lock}(((a, r, a), I, (0, []), \{\})) * \text{obs}(O)) \multimap Q(\text{tt}))) \vee$
as g_initc	$(\exists Wt, Ot, a, r, l, M, M'. \text{ucond}((a, r, a), (0, []), (0, []), \{\}) * \text{ulock}(l, Wt, Ot) * ((\text{cond}((a, r, A(l)), (0, []), M, M') * \text{ulock}(l, Wt, Ot)) \multimap Q(\text{tt}))) \vee$
as g_load	$(\exists v, l. (\text{cond}(v) * \text{ulock/locked}(l, Wt, Ot) * \text{obs}(O) \wedge L(v)=A(l)) * ((\text{cond}(v) * \text{ulock/locked}(l, Wt, Ot \uplus \{A(v)\}) * \text{obs}(O \uplus \{O(v)\})) \multimap Q(\text{tt})))$
as g_discharge	$(\exists v, l. (\text{cond}(v) * \text{ulock/locked}(l, Wt, Ot) * \text{obs}(O) \wedge \text{enoughObs}(v, Wt, Ot - \{A(v)\}) \wedge L(v)=A(l)) * ((\text{cond}(v) * \text{ulock/locked}(l, Wt, Ot - \{A(v)\}) * \text{obs}(O - \{O(v)\})) \multimap Q(\text{tt}))) \vee$
as g_new_ctr	$(\forall gv. \text{ctr}(gv, 0) \multimap Q(\text{tt})) \vee$
as g_inc	$(\exists n, gv. \text{ctr}(gv, n) * ((\text{ctr}(gv, n+1) * \text{tic}(gv)) \multimap Q(\text{tt}))) \vee$
as g_dec	$(\exists n, gv. \text{ctr}(gv, n) * \text{tic}(gv) * (\text{ctr}(gv, n-1) \multimap Q(\text{tt})))$

■ **Figure 42** Weakest precondition (part two of two).

weakest precondition of a command-context as shown in Definitions 17 and 18. Having these definitions, we can prove some auxiliary lemmas, shown in Lemmas 20, 21, 22, 23, and 24, which are used to prove Theorem 31.

► **Definition 17** (Weakest Precondition of a Context).

$$\text{wp}_n(\xi) = \begin{cases} \lambda _. \text{obs}(\{\}) & \text{if } \xi = \text{done} \\ \lambda z. \text{wp}_n(c[z/x], \text{wp}_n(\xi')) & \text{if } \xi = \text{let}'(x, c, \xi') \\ \lambda z. 0 < z ? \text{wp}_n(c_1, (\text{wp}_n(\xi'))) : \text{wp}_n(c_2, (\text{wp}_n(\xi'))) & \text{if } \xi = \text{if}'(c_1, c_2, \xi') \end{cases}$$

► **Definition 18** (Weakest Precondition of a command-context).

$$\text{wpc}_n(c, \xi) = \text{wp}_n(c, \text{wp}_n(\xi))$$

► **Lemma 19** (Weakening Postcondition).

$$p, \tilde{O}, g \models \text{wp}_n(c, Q) \wedge (\forall z. Q(z) \vdash Q'(z)) \Rightarrow \forall n' \leq n. p, \tilde{O}, g \models \text{wp}_{n'}(c, Q')$$

Proof. By induction on n and case analysis of c . ◀

► **Lemma 20** (Weakest Precondition of Wait).

$$\begin{aligned} \forall n, e_1, e_2, \xi, p, O, g. \quad & p, O, g \models \text{wpc}_n(\text{wait}(e_1, e_2), \xi) \Rightarrow \\ & \exists p_1, p_2, g_1, g_2, O_1, v, l, Wt, Ot. p = p_1 \uplus p_2 \wedge O = O_1 \uplus \{O(l)\} \wedge g = g_1 \uplus g_2 \wedge \\ & A(v) = \llbracket e_1 \rrbracket \wedge A(l) = \llbracket e_2 \rrbracket \wedge p_1(l) = \text{locked}(Wt, Ot) \wedge p_1(v) = \text{cond} \wedge \\ & p_2, \emptyset, g_2 \models l(l)(Wt \uplus \{A(v)\}, Ot) \wedge O(v) \prec O_1 \wedge O(l) \prec O_1 \uplus M'(v) \wedge L(v) = A(l) \wedge \\ & \text{enoughObs}(v, Wt \uplus \{A(v)\}, Ot) \wedge \\ & p_1[l := \text{lock}], O_1, g_1 \models \text{wpc}_n(\text{waiting4cvar}(e_1, e_2), \xi) \end{aligned}$$

► **Lemma 21** (Weakest Precondition of Notify).

$$\begin{aligned} \forall n, e, \xi, p, O, g. \quad & p, O, g \models \text{wpcx}_n(\text{notify}(e), \xi) \Rightarrow \exists p_1, p_M, g_1, g_M, O_1, v, l, Wt, Ot. \\ & p = p_1 \uplus p_M \wedge g = g_1 \uplus g_M \wedge O = O_1 \uplus (0 < Wt(A(v)) ? M'(v) : \{\}) \\ & \wedge A(v) = \llbracket e \rrbracket \wedge L(v) = A(l) \wedge p_1(v) = \text{cond} \wedge p_1(l) = \text{locked}(Wt, Ot) \wedge \\ & (0 < Wt(A(v)) ? p_M, \emptyset, g_M \models M(v) : (p_M = \mathbf{0} \wedge g_M = \mathbf{0})) \wedge \\ & p_1[l := \text{locked}(Wt - \{A(v)\}, Ot)], O_1, g_1 \models \text{wpcx}_n(\text{tt}, \xi) \end{aligned}$$

► **Lemma 22** (Weakest Precondition of waiting4cvar).

$$\begin{aligned} \forall n, e_1, e_2, \xi, p, O, g. \quad & p, O, g \models \text{wpcx}_n(\text{waiting4cvar}(e_1, e_2), \xi) \Rightarrow \\ & \exists v, l. p(v) = \text{cond} \wedge (p(l) = \text{lock} \vee \exists Wt, Ot. p(l) = \text{locked}(Wt, Ot)) \wedge L(v) = A(l) \wedge O(v) \prec O \\ & \wedge O(l) \prec O \uplus M'(v) \wedge A(v) = \llbracket e_1 \rrbracket \wedge A(l) = \llbracket e_2 \rrbracket \wedge \\ & \forall p_M, g_M. p_M, \emptyset, g_M \models M(v) \Rightarrow p \uplus p_M, O \uplus M'(v), g \uplus g_M \models \text{wpcx}_n(\text{waiting4lock}(e_2), \xi) \end{aligned}$$

► **Lemma 23** (Weakest Precondition of g_discharge).

$$\begin{aligned} \forall n, \xi, p, O, g. \quad & p, O, g \models \text{wpcx}_n(\text{g_discharge}, \xi) \Rightarrow \\ & \exists O_1, Wt, Ot, v, l. O = O_1 \uplus \{O(v)\} \wedge p(l) = \text{locked}(Wt, Ot) \wedge p(v) = \text{cond} \wedge \\ & \text{enoughObs}(v, Wt, Ot - \{A(v)\}) \wedge L(v) = A(l) \wedge \\ & p[l := \text{locked}(Wt, Ot - \{A(v)\})], O_1, g \models \text{wpcx}_n(\text{tt}, \xi) \end{aligned}$$

► **Lemma 24** (Weakest Precondition of fork).

$$\begin{aligned} \forall n, c, \xi, p, O, g. \quad & p, O, g \models \text{wpcx}_n(\text{fork}(c), \xi) \Rightarrow \\ & \exists p_1, p_2, g_1, g_2, O_1, O_2. p = p_1 \uplus p_2 \wedge g = g_1 \uplus g_2 \wedge O = O_1 \uplus O_2 \wedge \\ & p_1, O_1, g_1 \models \text{wpcx}_n(\text{tt}, \xi) \wedge p_2, O_2, g_2 \models \text{wp}_{n-1}(c, \lambda_.\text{obs}(\{\})) \end{aligned}$$

► **Lemma 25** (Frame in Weakest Precondition).

$$\forall n, c, Q, F, p, \tilde{O}, g. \quad p, \tilde{O}, g \models \text{wp}_n(c, Q) * F \Rightarrow \forall n' \leq n. p, \tilde{O}, g \models \text{wp}_{n'}(c, (\lambda z. Q(z) * F))$$

Proof. By induction on n and case analysis of c . ◀

D.4 Correctness of Commands

We define *correctness of commands*, as shown in Definition 26, ensuring that each proposed proof rule, where $\text{correct}_{pt}(P, c, Q)$ is abbreviated as $\{P\} c \{Q\}$, respects the definition of the weakest precondition. Having this definition we prove the proposed proof rules, ensuring deadlock freedom of importer channels, as well as some other necessary proof rules shown in Theorems 27, 28, and 29.

► **Definition 26** (Correctness of Commands). *A command is correct w.r.t a precondition P and a postcondition Q if and only if P implies the weakest precondition of that command w.r.t Q .*

$$\text{correct}_{pt}(P, c, Q) \Leftrightarrow \forall n. P \Rightarrow \text{wp}_{n,pt}(c, Q)$$

► **Theorem 27** (Rule Sequential Composition).

$$\text{correct}(P, c_1, Q) \wedge (\forall z. \text{correct}(Q(z), c_2[z/x], R)) \Rightarrow \text{correct}(P, \text{let}(x, c_1, c_2), R)$$

► **Theorem 28** (Rule Consequence).

$$\text{correct}(P, c, Q) \wedge (P' \vdash P) \wedge (\forall z. Q(z) \vdash Q'(z)) \Rightarrow \text{correct}(P', c, Q')$$

► **Theorem 29** (Rule Frame).

$$\text{correct}(P, c, Q) \Rightarrow \text{correct}(P * F, c, \lambda z. Q(z) * F)$$

As previously mentioned, since in this formalization ghost information is associated with lock and condition variable addresses via the lock and cond permission rather than via global function, we provide a new version of the proof rules, proposed in Section 3, regarding this formalization as shown in Figure 43.

D.5 Validity of a Configuration

We define *validity of a configuration*, shown in Definition 30, and prove that 1) starting from a valid configuration, all the subsequent configurations of the execution are also valid (Theorem 31), 2) a valid configuration is not deadlocked (Theorem 32), and 3) if a program c is verified by the proposed proof rules, where the verification starts from an empty bag of obligations and ends with such a bag too, then the initial configuration, where the heap is empty, denoted by $\mathbf{0} = \lambda _ . \emptyset$, and there is only one thread with the command c (and a context done), is a valid configuration (Theorem 34).

► **Definition 30** (Validity of a Configuration). *A configuration (t, h) is valid for n steps, denoted by $\text{valid}_n(t, h)$, if there exist a list of augmented threads A , consisting of the identification (id), the program (c), the context (ξ), the permission heap (p), the ghost resource heap (g) and the obligations (O) associated with each thread; a list of lock-invariant pairs Lin_v , storing the locks which are not held along with their invariants; three permission heaps p_i (associated with the invariants of the locks which are not held), p_l (the part of the permission heap which is leaked), and p_A (the union of all permission heaps in A and p_i as well as p_l); three ghost resource heaps g_i (associated with the invariants of the locks which are not held), g_l (the part of the ghost resource heap which is leaked), g_A (the union of all ghost resource heaps in A and g_i as well as g_l); and $locs$ (the set of locations for which a memory has been allocated), such that all of the following conditions hold:*

1. $\forall id, c, \xi. t(id) = (c; \xi) \Leftrightarrow \exists p, O, g. (id, c, \xi, p, O, g) \in A$
2. $\forall (id_1, c_1, \xi_1, p_1, O_1, g_1) \in A, (id_2, c_2, \xi_2, p_2, O_2, g_2) \in A. id_1 = id_2 \Rightarrow (id_1, c_1, \xi_1, p_1, O_1, g_1) = (id_2, c_2, \xi_2, p_2, O_2, g_2)$
3. $p_A = p_i \uplus p_l \uplus \biguplus_{(id, c, \xi, p, O, g) \in A} p \wedge g_A = g_i \uplus g_l \uplus \biguplus_{(id, c, \xi, p, O, g) \in A} g$
4. $\forall l_1, l_2. p_A(l_1) \neq \emptyset \wedge p_A(l_2) \neq \emptyset \wedge A(l_1) = A(l_2) \Rightarrow l_1 = l_2$
5. $\forall l. p_A(l) \neq \emptyset \Leftrightarrow l \in locs$
6. $\text{pheap_heap}(p_A, h)$
7. $p_i, \emptyset, g_i \models_{(l, inv) \in Lin_v}^* inv$
8. $p_A(l) = \text{lock} \wedge \neg \text{held}_h(l) \Rightarrow (l, I(l)(Wt_{l,A}, Ot_{l,A})) \in Lin_v$
9. $(l, inv) \in Lin_v \Rightarrow p_A(l) = \text{lock} \wedge \neg \text{held}_h(l)$
10. $\forall o \in O_A. \exists l. O(l) = o \wedge (p_A(l) = \text{cond} \vee p_A(l) = \text{lock} \vee \exists Wt, Ot. p_A(l) = \text{locked}(Wt, Ot))$
11. $p_A(l) = \text{unlock}(Wt, Ot) \vee p_A(l) = \text{locked}(Wt, Ot) \Rightarrow Wt = Wt_{l,A} \wedge Ot = Ot_{l,A}$
12. $p_A(l) = \text{lock} \vee p_A(l) = \text{unlock}(Wt, Ot) \vee p_A(l) = \text{locked}(Wt, Ot) \Rightarrow \text{held}_h(l) \Rightarrow l \in O_A$
13. $\forall (id, c, \xi, p, O, g) \in A.$
 - a. $p, O, g \models \text{wpcx}_n(c, \xi)$
 - b. $c = \text{waiting4cvar}(e_1, e_2) \Rightarrow \text{enoughObs}(\llbracket e_1 \rrbracket, Wt_{\llbracket e_2 \rrbracket, A}, Ot_{\llbracket e_2 \rrbracket, A})$

where

$$\blacksquare O_A = \biguplus_{(id, c, \xi, p, O, g) \in A} O$$

$$\begin{array}{l}
\text{NEWLOCK} \\
\{\text{true}\} \text{newlock } \{\lambda a. \text{unlock}(((a, r, a), (0, []), (0, []), \{\}, \{\}, \{\})))\} \\
\\
\text{INITLOCK} \\
\{\text{unlock}(((a, r, a), (0, []), (0, []), \{\}), Wt, Ot) * pt(I_{index}, I_{args})(Wt, Ot) * \text{obs}(O)\} \text{nop} \\
\{\lambda _ . \text{lock}(((a, r, a), (I_{index}, I_{args}), (0, []), \{\})) * \text{obs}(O)\} \\
\\
\text{ACQUIRE} \\
\{\text{lock}(l) * \text{obs}(O) \wedge O(l) \prec O \wedge A(l) = a_l\} \text{acquire}(a_l) \\
\{\lambda _ . \exists Wt, Ot. \text{locked}(l, Wt, Ot) * l(l)(Wt, Ot) * \text{obs}(O \uplus \{O(l)\})\} \\
\\
\text{RELEASE} \\
\{\text{locked}(l, Wt, Ot) * l(l)(Wt, Ot) * \text{obs}(O \uplus \{O(l)\}) \wedge A(l) = a_l\} \text{release}(a_l) \\
\{\lambda _ . \text{lock}(l) * \text{obs}(O)\} \\
\\
\text{NEWCV} \\
\{\text{true}\} \text{new_cvar } \{\lambda a. \text{ucond}(((a, r, a), (0, []), (0, []), \{\})))\} \\
\\
\text{INITCV} \\
\{\text{ucond}((a, r, a), (0, []), (0, []), \{\}) * \text{unlock}(l, Wt, Ot)\} \text{nop} \\
\{\lambda _ . \text{cond}((a, r, A(l)), (0, []), (M_{index}, M_{args}), M') * \text{unlock}(l, Wt, Ot)\} \\
\\
\text{WAIT} \\
\{\text{cond}(v) * \text{locked}(l, Wt, Ot) * l(l)(Wt \uplus \{A(v)\}, Ot) * \text{obs}(O \uplus \{O(l)\}) \wedge A(v) = a_v \wedge A(l) = a_l \\
\wedge A(l) = L(v) \wedge O(v) \prec O \wedge O(l) \prec O \uplus M'(v) \wedge \text{enoughObs}(A(v), Wt \uplus \{A(v)\}, Ot)\} \text{wait}(a_v, a_l) \\
\{\lambda _ . \text{cond}(v) * \text{obs}(O \uplus \{O(l)\}) \uplus M'(v) * \exists Wt', Ot'. \text{locked}(l, Wt', Ot') * l(l)(Wt', Ot') * M(v)\} \\
\\
\text{NOTIFY} \\
\{\text{obs}(O \uplus (0 < Wt(A(v)) ? M'(v) : \{\})) * \text{cond}(v) * \text{locked}(l, Wt, Ot) * (Wt(A(v)) = 0 \vee M(v)) \wedge \\
A(l) = L(v) \wedge A(v) = a_v\} \text{notify}(a_v) \{\lambda _ . \text{obs}(O) * \text{cond}(v) * \text{locked}(l, Wt - \{A(v)\}, Ot)\} \\
\\
\text{NOTIFYALL} \\
\{\text{cond}(v) * \text{locked}(l, Wt, Ot) * (\prod_{i=1}^{Wt(A(v))} M(v)) \wedge M'(v) = \{\} \wedge A(l) = L(v) \wedge A(v) = a_v\} \\
\text{notifyAll}(a_v) \{\lambda _ . \text{cond}(v) * \text{locked}(l, Wt[v:=0], Ot)\} \\
\\
\text{CHARGE OBLIGATION} \\
\{\text{obs}(O) * \text{cond}(v) * \text{unlock/locked}(l, Wt, Ot) \wedge A(l) = L(v)\} \text{nop} \\
\{\lambda _ . \text{obs}(O \uplus \{O(v)\}) * \text{cond}(v) * \text{unlock/locked}(l, Wt, Ot \uplus \{A(v)\})\} \\
\\
\text{DISCHARGE OBLIGATION} \\
\{\text{obs}(O) * \text{cond}(v) * \text{unlock/locked}(l, Wt, Ot) \wedge \text{enoughObs}(A(v), Wt, Ot - \{A(v)\}) \\
\wedge A(l) = L(v)\} \text{nop} \{\lambda _ . \text{obs}(O - \{O(v)\}) * \text{unlock/locked}(l, Wt, Ot - \{A(v)\})\}
\end{array}$$

■ **Figure 43** Proof rules verifying deadlock-freedom of importer monitors, where ghost information is associated with lock and channel addresses via lock and cond permissions.

22:56 Transferring Obligations Through Synchronizations

- $O_{t,A} = \lambda v. L(v)=l ? O_A(v) : 0$, and $Wt_{l,A} = \biguplus_{(id,c,\xi,p,O,g) \in A \wedge \text{waiting_for}_h(c)=v \wedge L(v)=l} \{v\}$
- $\text{waiting_for}_h(c)$ returns the object for which c is waiting, if any, i.e.

$$\text{waiting_for}_h(c, h) = \begin{cases} \llbracket e_1 \rrbracket & \text{if } c = \text{waiting4cond}(e_1, e_2) \\ \llbracket e \rrbracket & \text{if } c = \text{waiting4lock}(e) \wedge h(\llbracket e \rrbracket) \neq 1 \\ \emptyset & \text{otherwise} \end{cases}$$
- $\text{held}_h(l)$ returns true if and only if the lock l is held, i.e. $\text{held}_h(l) \Rightarrow h(A(l)) \neq 1$

Each item in this definition ensures some properties as follows: (1) and (2) ensure that the list of augmented threads A is correctly associated with the thread tables t , (3) ensures that the union of all permission heaps as well as the union of all ghost resource heaps in A are defined, (4) ensures that any two allocated locations which have the same address are equal, (5) $locs$ is the set of locations for which a memory has been allocated (having this set it is possible to map addresses to their ghost information) (6) ensures that p_A corresponds to the concrete heap h , (7) ensures that p_i and g_i model the separating conjunction of the invariants of the locks which are not held, and these invariants do not assert any obligation, (8) ensures that any lock which is not held along with its invariant exist in Linv , (9) ensures that the locks in Linv are not held, (10) ensures that any obligation in A is associated with the address of a condition variable or a lock which has been initialized, (11) ensures that the parameters Wt and Ot stored in the permissions `unlock` and `locked` of any lock store the total number of waiting threads and obligations of the condition variables associated with that lock, respectively, (12) ensures that any lock l which is held exists in the union of the obligations of A , (13-a) the permission heap, the obligations, and the ghost resource heap of each thread model the weakest precondition of the command of that thread w.r.t. the postcondition in which there is no obligation, and (13-b) for any condition variable for which a thread is waiting the invariant `enoughObs` holds.

► **Theorem 31** (Small Steps Preserve Validity of Configurations). *Each step of the execution preserves validity of configurations.*

$$(t, h) \rightsquigarrow (t', h') \wedge \text{valid}_{n+1}(t, h) \Rightarrow \text{valid}_n(t', h')$$

Proof. By case analysis of the small step relation \rightsquigarrow . ◀

► **Theorem 32** (A Valid Configuration Is Not Deadlocked). *If a valid configuration has some threads then there exists a thread in this configuration neither waiting for a condition variable nor a lock.*

$$\text{valid}_n(t, h) \wedge \exists id. t(id) \neq \emptyset \Rightarrow \exists id'. \text{waiting_for}(\text{fst}(t(id'))) = \emptyset$$

Proof. We assume that all threads in t are waiting for a condition variable or a lock. Since (t, h) is a valid configuration there exists a valid list of augmented threads A with a corresponding valid bag $G = \text{valid_bag}(A)$, where `valid_bag` maps any element (id, c, ξ, p, O, g) to an element $(\text{waiting_for}(c), O)$. By Lemma 33, we have $G = \emptyset$, implying $A = \{\}$, implying $t = \mathbf{0}$ which contradicts the hypothesis of the theorem.

Note that in the definition of validity of a configuration we also keep track of all locations whose addresses are allocated, which makes it possible to provide the function R , mapping lock and condition variable addresses to their levels, for Lemma 33. The first hypothesis in this lemma is met by the constraint 13-a in the definition of validity of a configuration. Additionally, the second hypothesis in this lemma is met by the constraints 12 or 13-b, where the related thread is waiting for a lock or a condition variable, respectively. ◀

$$\begin{aligned}
\text{ol}(m, r) &::= (m.l, r, m.l) \\
\text{ov}(m, r) &::= (m.v, r, m.l) \\
l(m, r) &::= (\text{ol}(m, r), (\text{linv}, [m]), (0, []), \{\}) \\
v(m, r) &::= (\text{ov}(m, r), (0, []), (\mathbf{M}, []), \{\text{ov}(m, r)\}) \\
\text{mutex}(\text{mutex } m, \text{waitobj } o) &= \text{lock}(l(m, R(o)-1)) * \text{cond}(o) \wedge o = v(m, R(o)) \\
\text{pt}(\text{linv}, [m].\text{args}) &= \lambda Wt. \lambda Ot. \exists b, w. m.b \mapsto b * m.w \mapsto w \wedge Wt(m.v)=w \wedge \\
&\quad (0 < b ? 0 < Ot(v) : Wt(v) = 0) \\
\text{pt}(\mathbf{M}, \text{args}) &= \lambda Wt. \lambda Ot. \text{true}
\end{aligned}$$

■ **Figure 44** Verification of the fair mutexes implementation shown in Figure 16 using the proof rules in Figure 43 (part one of two).

Lemma 33 ensures that in any state of the execution if all the desired invariants are respected then it is impossible that all threads are waiting for an object. In this lemma G is a bag of object-obligations pairs such that each element t of G is associated with a thread in a state of the execution, where the first element of t is associated with the object for which t is waiting and the second element is associated with the obligations of t .

► **Lemma 33** (A Valid Bag of Augmented Threads Is Not Deadlocked).

$$\begin{aligned}
&\forall G : \text{Bags}(\text{Addresses} \times \text{Bags}(\text{Addresses})), R : \text{Addresses} \rightarrow \text{Levels}. \\
&(\forall (o, O) \in G. o \prec O \wedge (\exists o', O'. (o', \{o\} \uplus O') \in G)) \Rightarrow G = \{\}
\end{aligned}$$

Proof. By contradiction; assume that $\exists (o_m, O_1) \in G$ where $\neg \exists (o, O) \in G. R(o) < R(o_m)$. By H_2 we have $\exists o', O'. (o', \{o_m\} \uplus O') \in G$ and by H_1 we have $o' \prec \{o_m\} \uplus O'$, which contradicts minimality of the level of o_m . ◀

► **Theorem 34** (The Initial Configuration Is Valid). *The initial configuration, consisting of an empty heap and a single thread whose program is verified by the proposed proof rules, is a valid configuration.*

$$\text{correct}_{sp}(\text{obs}(\{\}), c, \lambda _ . \text{obs}(\{\})) \Rightarrow \forall n, id. \text{valid}_n(\mathbf{0}[id:=c; \text{done}], \mathbf{0})$$

Proof. The goal is achieved because there are an augmented thread list $T = [(id, c, \text{done}, \mathbf{0}, \{\}, \mathbf{0})]$, a list of lock-invariant pairs $\text{Linv} = []$, two permission heaps $p_i = \mathbf{0}$ and $p_l = \mathbf{0}$, and two ghost resource heaps $g_i = \mathbf{0}$ and $g_l = \mathbf{0}$, such that all the conditions in the definition of validity of configurations are met. ◀

D.6 An Example Proof

In this section we show how the program in Figure 16 can be verified using the proof rules in Figure 43, as shown in Figures 44 and 45¹⁶.

¹⁶Note that for this program we assume a straightforward extension of the programming language with immutable structures, i.e. tuples with named components.

```

routine new_mutex(){
req : {true}
  l := new_lock;
  {unlock(((l, r-1, l), (0, []), (0, []), {}, {}), {}, {})}
  v := new_cvar; nop; //Rule INITCV
  {unlock(((l, r-1, l), (0, []), (0, []), {}, {}), {}, {}) * cond((v, r, l), (0, []), (M, []), {(v, r, l)})}
  m := mutex(l:=l, v:=v, b:=new_int(1), w:=new_int(1)); nop; //Rule INITLOCK
  m ens : {λm. ∃o. mutex(m, o) ∧ R(o)=r}

routine enter_cs(mutex m){
req : {obs(O) * mutex(m, o) ∧ O(o) < O}
  acquire(m.l);
  // Let l = l(m, R(o)-1) and ol = ol(m, R(o)-1)
  {obs(O⊕{ol}) * mutex(m, o) * ∃Wt, Ot. locked(l, Wt, Ot) * pt(linv, [m])(Wt, Ot)}
  if(m.b){
    m.w := m.w+1;
    {obs(O⊕{ol}) * mutex(m, o) * ∃Wt, Ot. locked(l, Wt, Ot) * pt(linv, [m])(Wt⊕{m.v}, Ot)}
    wait(m.v, m.l)
    {obs(O⊕{ol, O(o)}) * mutex(m, o) * ∃Wt, Ot. locked(l, Wt, Ot) * pt(linv, [m])(Wt, Ot)}
  } else { m.b := 1; nop; //Rule CHARGE OBLIGATION
    {obs(O⊕{ol, O(o)}) * mutex(m, o) * ∃Wt, Ot. locked(l, Wt, Ot) * pt(linv, [m])(Wt, Ot)};
    {obs(O⊕{ol, O(o)}) * mutex(m, o) * ∃Wt, Ot. locked(l, Wt, Ot) * pt(linv, [m])(Wt, Ot)}
    release(m.l)
  }
  m ens : {obs(O⊕{O(o)}) * mutex(m, o)}

routine exit_cs(mutex m){
req : {obs(O⊕{O(o)}) * mutex(m, o) ∧ O(o) < O}
  acquire(m.l);
  // Let l = l(m, R(o)-1) and ol = ol(m, R(o)-1)
  {obs(O⊕{ol, O(o)}) * mutex(m, o) * ∃Wt, Ot. locked(l, Wt, Ot) * pt(linv, [m])(Wt, Ot)}
  m.b := 0;
  if(0 < m.w){ m.w := m.w-1; m.b := 1; notify(m.v)
    {obs(O⊕{ol}) * mutex(m, o) * ∃Wt, Ot. locked(l, Wt, Ot) * pt(linv, [m])(Wt, Ot)}
  } else { nop //Rule DISCHARGE OBLIGATION
    {obs(O⊕{ol}) * mutex(m, o) * ∃Wt, Ot. locked(l, Wt, Ot) * pt(linv, [m])(Wt, Ot)};
    release(m.l)
  }
  m ens : {obs(O) * mutex(m, o)}

```

■ **Figure 45** Verification of the fair mutexes implementation shown in Figure 16 using the proof rules in Figure 43 (part two of two).