

Derivation of Heard-Of Predicates From Elementary Behavioral Patterns

Adam Shimi, Aurélie Hurault, and Philippe Queinnec

IRIT – Université de Toulouse, 2 rue Camichel, F-31000 Toulouse, France
`{firstname}.{lastname}@irit.fr`

Abstract. There are many models of distributed computing, and no unifying mathematical framework for considering them all. One way to sidestep this issue is to start with simple communication and fault models, and use them as building blocks to derive the complex models studied in the field. We thus define operations like union, succession or repetition, which makes it easier to build complex models from simple ones while retaining expressivity.

To formalize this approach, we abstract away the complex models and operations in the Heard-Of model. This model relies on (possibly asynchronous) rounds; sequence of digraphs, one for each round, capture which messages sent at a given round are received before the receiver goes to the next round. A set of sequences, called a heard-of predicate, defines the legal communication behaviors – that is to say, a model of communication. Because the proposed operations behave well with this transformation of operational models into heard-of predicates, we can derive bounds, characterizations, and implementations of the heard-of predicates for the constructions.

Keywords: Message-passing · Asynchronous Rounds · Failures · Heard-Of Model

1 Introduction

1.1 Motivation

Let us start with a round-based distributed algorithm; such an algorithm is quite common in the literature, especially in fault-tolerant settings. We want to formally verify this algorithm using the methods of our choice: proof-assistant, model-checking, inductive invariants, abstract interpretation. . . But how are we supposed to model the context in which the algorithm will run? Even a passing glance at the distributed computing literature shows a plethora of models defined in the mixture of english and mathematics.

Thankfully, there are formalisms for abstracting round-based models of distributed computing. One of these is the Heard-Of model of Charron-Bost and Schiper [4]; it boils down the communication model to a description of all accepted combinations of received messages. Formally, this is done by considering

communications graphs, one for each round, and taking the sets of infinite sequences of graphs that are allowed by the model. Such a set is called a heard-of predicate, and captures a communication model.

An angle of attack for verification is therefore to find the heard-of predicate corresponding to a real-world environment, and use the techniques from the literature to verify an algorithm for this heard-of predicate. But which heard-of predicate should be used? What is the "right" predicate for a given environment? For some cases, the predicates are given in Charron-Bost and Schiper [4]; but this does not solve the general case.

Actually, the answer is quite subtle. This follows from a fundamental part of the Heard-Of model: communication-closedness [7]. This means that for p to use a message from q at round r , p must receive it before or during its own round r . And thus, knowing whether p receives the message from q at the right round or not depends on how p waits for messages. That is, it depends on the specifics of how rounds are implemented on top of it.

Once again, the literature offers a solution: Shimi et al. [12] propose to first find a delivered predicate – a description of which messages will eventually be delivered, without caring about rounds –, and then to derive the heard-of predicate from it. This derivation explicitly studies strategies, the aforementioned rules for how processes wait for messages before changing round.

But this brings us back to square one: now we are looking for the delivered predicate corresponding to a real-world model, instead of the heard-of predicate. Basic delivered predicates for elementary failures are easy to find, but delivered predicates corresponding to combinations of failures are often not intuitive.

In this paper, we propose a solution to this problem: building a complex delivered predicate from simpler ones we already know. For example, consider a system where one process can crash and may recover later, and another process can definitively crash. The delivered predicate for at most one crash is $PDel_1^{crash}$, and the predicate where all the messages are delivered is $PDel^{total}$. Intuitively, a process that can crash and necessarily recover is described by the behavior of $PDel_1^{crash}$ followed by the behavior of $PDel^{total}$. We call this the succession of these predicates, and write it $PDel_1^{recover} \triangleq PDel_1^{crash} \rightsquigarrow PDel^{total}$. In our system, the crashed process may never recover: hence we have either the behavior of $PDel_1^{recover}$ or the behavior of $PDel_1^{crash}$. This amounts to a union (or a disjunction); we write it $PDel_1^{canrecover} \triangleq PDel_1^{recover} \cup PDel_1^{crash}$. Finally, we consider a potential irremediable crash, additionally to the previous predicate. Thus we want the behavior of $PDel_1^{crash}$ and the behavior of $PDel_1^{canrecover}$. We call it the combination (or conjunction) of these predicates, and write it $PDel_1^{crash} \otimes PDel_1^{canrecover}$. The complete system is thus described by $PDel_1^{crash} \otimes ((PDel_1^{crash} \rightsquigarrow PDel^{total}) \cup PDel_1^{crash})$. In the following, we will also introduce an operator ω to express repetition. For example, a system where, repeatedly, a process can crash and recover is $(PDel_1^{crash} \rightsquigarrow PDel^{total})^\omega$.

Lastly, the analysis of the resulting delivered predicate can be bypassed: its heard-of predicate arises from our operations applied to the heard-of predicates of the elementary building blocks.

1.2 Related Work

The heard-of model was proposed by Charron-Bost and Schiper [4] as a combination of the ideas of two previous work. First, the concept of a fault model where the only information is which message arrives, from Santoro and Widmayer [11]; and second, the idea of abstracting failures in a round per round fashion, from Gafni [8]. Replacing the operational fault detectors of Gafni with the fault model of Santoro and Widmayer gives the heard-of model.

This model was put to use in many ways. Obviously computability and complexity results were proven: new algorithms for consensus in the original paper by Charron-Bost and Schiper [4]; characterizations for consensus solvability by Coulouma et al. [5] and Nowak et al. [10]; a characterization for approximate consensus solvability by Charron-Bost et al. [3]; a study of k set-agreement by Biely et al. [1]; and more.

The clean mathematical abstraction of the heard-of model also works well with formal verification. The rounds provide structure, and the reasoning can be less operational than in many distributed computing abstractions. We thus have a proof assistant verification of consensus algorithms in Charron-Bost et al. [2]; cutoff bounds for the model checking of consensus algorithms by Marić et al. [9]; a DSL to write code following the structure of the heard-of model and verify it with inductive invariants by Drăgoi et al. [6]; and more.

1.3 Contributions

The contributions of the paper are:

- A definition of operations on delivered predicates and strategies, as well as examples using them in Section 2.
- The study of oblivious strategies, the strategies only looking at messages for the current round, in Section 3. We provide a technique to extract a strategy dominating the oblivious strategies of the built predicate from the strategies of the initial predicates; exact computations of the generated heard-of predicates; and a sufficient condition on the building blocs for the result of operations to be dominated by an oblivious strategy.
- The study of conservative strategies, the strategies looking at everything but messages from future rounds, in Section 4. We provide a technique to extract a strategy dominating the conservative strategies of the build predicate from the strategies of the initial predicates; upper bounds on the generated heard-of predicates; and a sufficient condition on the building blocs for the result of operations to be dominated by a conservative strategy.

Due to size constraints, many of the complete proofs are not in the paper itself, and can be found in the appendix.

2 Operations and Examples

2.1 Basic concepts

We start by providing basic definitions and intuitions. The model we consider proceed by rounds, where processes send messages tagged with a round number, wait for some messages with this round number, and then compute the next state and increment the round number.

Definition 1 (Collections and Predicates). *Let Π a set of processes. Then an element of $(\mathbb{N}^* \times \Pi) \mapsto \mathcal{P}(\Pi)$ is either a **Delivered collection** c or a **Heard-Of collection** h for Π , depending on the context. c_{tot} is the total collection such that $\forall r > 0, \forall p \in \Pi : c_{tot}(r, p) = \Pi$.*

An element of $\mathcal{P}((\mathbb{N}^ \times \Pi) \mapsto \mathcal{P}(\Pi))$ is either a **Delivered predicate** $PDel$ or a **Heard-Of predicate** PHO for Π . $\mathcal{P}_{tot} = \{c_{tot}\}$ is the total delivered predicate.*

For a heard-of collection h , $h(r, p)$ are the senders of messages for round r that p has received at or before its round r , and thus has known while at round r . For a delivered collection c , $c(r, p)$ are the senders of messages for round r that p has received, at any point in time. Some of these messages may have arrived early, before p was at r , or too late, after p has left round r . c gives an operational point of view (which messages arrive), and h gives a logical point of view (which messages are used).

Remark 1. We also regularly use the "graph-sequence" notation for a collection c . Let $Graphs_{\Pi}$ be the set of graphs whose nodes are the elements of Π . Then a collection gr is an element of $(Graphs_{\Pi})^{\omega}$. We say that c and gr represent the same collection when $\forall r > 0, \forall p \in \Pi : c(r, p) = In_{gr[r]}(p)$, where $In(p)$ is the incoming vertices of p . We will usually not define two collections but use one collection as both kind of objects; the actual type being used in a particular expression can be deduced from the operations on the collection. For example $c[r]$ makes sense for a sequence of graphs, while $c(r, p)$ makes sense for a function.

In an execution, the local state of a process is the pair of its current round and all the received messages up to this point. A message is represented by a pair $\langle round, sender \rangle$. For a state q , $q(i)$ is the set of peers from which the process has received a message for round i .

Definition 2 (Local State). *Let $Q = \mathbb{N}^* \times \mathcal{P}(\mathbb{N}^* \times \Pi)$. Then $q \in Q$ is a **local state**.*

For $q = \langle r, mes \rangle$, we write $q.round$ for r , $q.mes$ for mes and $\forall i > 0 : q(i) \triangleq \{k \in \Pi \mid \langle i, k \rangle \in q.mes\}$.

We then define strategies, which constrain the behavior of processes. A strategy is a set of states from which a process is allowed to change round. It captures rules like "wait for at least F messages from the current round", or "wait for these specific messages". Strategies give a mean to constrain executions.

Definition 3 (Strategy). *$f \in \mathcal{P}(Q)$ is a **strategy**.*

2.2 Definition of Operations

We can now define operations on predicates and their corresponding strategies.

Definition 4 (Operations on predicates). *Let P_1, P_2 be two delivered or heard-of predicates.*

- The **union** of P_1 and P_2 is $P_1 \cup P_2$.
- The **combination** $P_1 \otimes P_2 \triangleq \{c_1 \otimes c_2 \mid c_1 \in P_1, c_2 \in P_2\}$, where for c_1 and c_2 two collections, $\forall r > 0, \forall p \in \Pi : (c_1 \otimes c_2)(r, p) = c_1(r, p) \cap c_2(r, p)$.
- The **succession** $P_1 \rightsquigarrow P_2 \triangleq \bigcup_{c_1 \in P_1, c_2 \in P_2} c_1 \rightsquigarrow c_2$,
with $c_1 \rightsquigarrow c_2 \triangleq \{c \mid \exists r \geq 0 : c = c_1[1, r].c_2\}$.
- The **repetition** of P_1 , $(P_1)^\omega \triangleq \{c \mid \exists (c_i)_{i \in \mathbb{N}^*}, \exists (r_i)_{i \in \mathbb{N}^*} : r_1 = 0 \wedge \forall i \in \mathbb{N}^* : (c_i \in P_1 \wedge r_i < r_{i+1} \wedge c[r_i + 1, r_{i+1}] = c_i[1, r_{i+1} - r_i])\}$.

The intuition behind these operations is the following:

- The union of two delivered predicates is equivalent to an OR on the two communication behaviors. For example, the union of the delivered predicate for one crash at round r and of the one for one crash at round $r + 1$ gives a predicate where there is either a crash at round r or a crash at round $r + 1$.
- The combination of two behaviors takes every pair of collections, one from each predicate, and computes the intersection of the graphs at each round. Meaning, it adds the loss of messages from both, to get both behaviors at once. For example, combining $PDel_1^{crash}$ with itself gives $PDel_2^{crash}$, the predicate with at most two crashes.
- For succession, the system starts with one behavior, then switch to another. The definition is such that the first behavior might never happen, but the second one must appear.
- Repetition is the next logical step after succession: instead of following one behavior with another, the same behavior is repeated again and again. For example, taking the repetition of at most one crash results in a potential infinite number of crash-and-restart, with the constraint of having at most one crashed process at any time.

For all operations on predicates, we provide an analogous one for strategies. We show later that strategies for the delivered predicates, when combined by the analogous operation, retain important properties on the result of the operation on the predicates.

Definition 5 (Operations on strategies). *Let f_1, f_2 be two strategies.*

- Their **union** $f_1 \cup f_2 \triangleq$ the strategy such that $\forall q$ a local state: $(f_1 \cup f_2)(q) \triangleq f_1(q) \vee f_2(q)$.
- Their **combination** $f_1 \otimes f_2 \triangleq \{q_1 \otimes q_2 \mid q_1 \in f_1 \wedge q_2 \in f_2 \wedge q_1.\text{round} = q_2.\text{round}\}$, where for q_1 and q_2 at the same round r , $q_1 \otimes q_2 \triangleq \langle r \{ \langle r', k \rangle \mid r' > 0 \wedge k \in q_1(r') \cap q_2(r') \} \rangle$

- Their **succession** $f_1 \rightsquigarrow f_2 \triangleq f_1 \cup f_2 \cup \{q_1 \rightsquigarrow q_2 \mid q_1 \in f_1 \wedge q_2 \in f_2\}$ where
 $q_1 \rightsquigarrow q_2 \triangleq \langle q_1.\text{round} + q_2.\text{round}, \{\langle r, k \rangle \mid r > 0 \wedge \begin{cases} k \in q_1(r) & \text{if } r \leq q_1.\text{round} \\ k \in q_2(r - q_1.\text{round}) & \text{if } r > q_1.\text{round} \end{cases} \rangle \rangle$
- The **repetition** of f_1 , $f_1^\omega \triangleq \{q_1 \rightsquigarrow q_2 \rightsquigarrow \dots \rightsquigarrow q_k \mid k \geq 1 \wedge q_1, q_2, \dots, q_k \in f_1\}$.

The goal is to derive new strategies for the resulting model by applying operations on strategies for the starting models. This allows, in some cases, to bypass strategies, and deduce the Heard-Of predicate for a given Delivered predicate from the Heard-Of predicates of its building blocks.

2.3 Executions and Domination

Before manipulating predicates and strategies, we need to define what is an execution: a specific ordering of events corresponding to a delivered collection.

Definition 6 (Execution). Let Π be a set of n processes. Let the set of transitions $T = \{\text{next}_j \mid j \in \Pi\} \cup \{\text{deliver}(r, k, j) \mid r \in \mathbb{N}^* \wedge k, j \in \Pi\} \cup \{\text{stop}\}$. next_j is the transition for j changing round, $\text{deliver}(r, k, j)$ is the transition for the delivery to j of the message sent by k in round r , stop models a deadlock. Then, $t \in T^\omega$ is an **execution** \triangleq

- (**Delivery after sending**)
 $\forall i \in \mathbb{N} : t[i] = \text{deliver}(r, k, j) \implies |\{l \in [0, i] \mid t[l] = \text{next}_k\}| \geq r - 1$
- (**Unique delivery**)
 $\forall \langle r, k, j \rangle \in (\mathbb{N}^* \times \Pi \times \Pi) : |\{i \in \mathbb{N} \mid t[i] = \text{deliver}(r, k, j)\}| \leq 1$
- (**Once stopped, forever stopped**)
 $\forall i \in \mathbb{N} : t[i] = \text{stop} \implies \forall j \geq i : t[j] = \text{stop}$

Let c be a delivered collection. Then, $\text{execs}(c)$, the **executions of c** \triangleq

$$\left\{ t \text{ an execution} \left| \begin{array}{l} \forall \langle r, k, j \rangle \in \mathbb{N}^* \times \Pi \times \Pi : \\ (k \in c(r, j) \wedge |\{i \in \mathbb{N} \mid t[i] = \text{next}_k\}| \geq r - 1) \\ \iff \\ (\exists i \in \mathbb{N} : t[i] = \text{deliver}(r, k, j)) \end{array} \right. \right\}$$

For $P\text{Del}$ a delivered predicate, we write
 $\text{execs}(P\text{Del}) = \{\text{execs}(c) \mid c \in P\text{Del}\}$.

Let t be an execution, $p \in \Pi$ and $i \in \mathbb{N}$. Then the state of p in t after i transitions is $q_p^t[i] \triangleq \langle |\{l < i \mid t[l] = \text{next}_p\}| + 1, \{\langle r, k \rangle \mid \exists l < i : t[l] = \text{deliver}(r, k, p)\} \rangle$

Recall that strategies constrain when processes can change round. Thus, the executions that conform to a strategy change rounds only when allowed by it, and do it infinitely often if possible.

Definition 7 (Executions of a Strategy). Let f be a strategy and t an execution. Then t is an **execution of f** $\triangleq t$ satisfies:

- (**All nexts allowed**) $\forall i \in \mathbb{N}, \forall p \in \Pi : (t[i] = \text{next}_p \implies q_p^t[i] \in f)$

- (**Fairness**) $\forall p \in \Pi : |\{i \in \mathbb{N} \mid t[i] = next_p\}| < \aleph_0 \implies |\{i \in \mathbb{N} \mid q_p^t[i] \notin f\}| = \aleph_0$

For a delivered predicate $PDel$, we note $execs_f(PDel) \triangleq \{t \mid t \text{ an execution of } f \wedge t \in execs(PDel)\}$.

An important part of this definition considers executions where processes cannot necessarily change round after each delivery. That is, in the case of "waiting for at most F messages", an execution where more messages are delivered than F at some round is still an execution of the strategy. This hypothesis captures the asynchrony of processes, which are not always scheduled right after deliveries. It is compensated by a strong fairness assumption: if a strategy allows infinitely often the change of round, it must eventually happen.

Going back to strategies, not all of them are worth using. Notably, strategies that block forever at some round are worthless. The validity of a strategy captures the absence of such an infinite wait.

Definition 8 (Validity).

An execution t is **valid** $\triangleq \forall p \in \Pi : |\{i \in \mathbb{N} \mid t[i] = next_p\}| = \aleph_0$.

Let $PDel$ a delivered predicate and f a strategy. f is a **valid strategy** for $PDel$ $\triangleq \forall t \in execs_f(PDel) : t$ is a valid execution.

Because in a valid execution no process is ever blocked at a given round, there are infinitely many rounds. Hence, the messages delivered before the changes of round uniquely define a heard-of collection.

Definition 9 (Heard-Of Collection of Executions and Heard-Of Predicate of Strategies). Let t be a valid execution. h_t is the **heard-of collection** of t \triangleq

$$\forall r \in \mathbb{N}^*, \forall p \in \Pi : h_t(r, p) = \left\{ k \in \Pi \mid \exists i \in \mathbb{N} : \begin{pmatrix} q_p^t[i].round = r \\ \wedge t[i] = next_p \\ \wedge \langle r, k \rangle \in q_p^t[i].mes \end{pmatrix} \right\}$$

Let $PDel$ be a delivered predicate, and f be a valid strategy for $PDel$. Then we write $PHO_f(PDel)$ for the heard-of predicate composed of the collections of the executions of f on $PDel$: $PHO_f(PDel) \triangleq \{h_t \mid t \in execs_f(PDel)\}$.

Lastly, the heard-of predicate of most interest is the smallest one that can be generated by a valid strategy on the delivered predicate. The intuition boils down to two ideas: less collections means more constrained communication, with more constrained communication meaning more powerful models; and if a heard-of predicate is contained in all the other predicates generated on this $PDel$, it implies them, and thus characterizes all the predicates that can be generated on this predicate. This notion of smallest predicate is formalized through an order on strategies and their heard-of predicates.

Definition 10 (Domination). Let $PDel$ be a delivered predicate and let f and f' be two valid strategies for $PDel$. Then f **dominates** f' for $PDel$, written $f' \prec_{PDel} f$, $\triangleq PHO_{f'}(PDel) \supseteq PHO_f(PDel)$.

A greatest element for \prec_{PDel} is called a **dominating strategy** for $PDel$. Given such a strategy f , the **dominating predicate** for $PDel$ is $PHO_f(PDel)$.

Description	Expression	HO	Proof
At most 1 crash	$\mathcal{P}_1^{crash} = \bigcup_{i=1}^{\infty} \mathcal{P}_{1,i}^{crash}$	$HOProd(\{T \subseteq \Pi \mid T \geq n-1\})$	[12, Thm 15]
At most F crashes	$\mathcal{P}_F^{crash} = \bigotimes_{j=1}^F \mathcal{P}_1^{crash}$	$HOProd(\{T \subseteq \Pi \mid T \geq n-F\})$	[12, Thm 15]
At most 1 crash, which will restart	$\mathcal{P}_1^{recover} = \mathcal{P}_1^{crash} \rightsquigarrow \mathcal{P}^{total}$	$HOProd(\{T \subseteq \Pi \mid T \geq n-1\})$	Thm 4
At most F crashes, which will restart	$\mathcal{P}_F^{recover} = \bigotimes_{j=1}^F \mathcal{P}_1^{recover}$	$HOProd(\{T \subseteq \Pi \mid T \geq n-F\})$	Thm 4
At most 1 crash, which can restart	$\mathcal{P}_1^{canrecover} = \mathcal{P}_1^{recover} \cup \mathcal{P}_1^{crash}$	$HOProd(\{T \subseteq \Pi \mid T \geq n-1\})$	Thm 4
At most F crashes, which can restart	$\mathcal{P}_F^{canrecover} = \bigotimes_{j=1}^F \mathcal{P}_1^{canrecover}$	$HOProd(\{T \subseteq \Pi \mid T \geq n-F\})$	Thm 4
No bound on crashes and restart, with only 1 crash at a time	$\mathcal{P}_1^{recovery} = (\mathcal{P}_1^{crash})^\omega$	$HOProd(\{T \subseteq \Pi \mid T \geq n-1\})$	Thm 4
No bound on crashes and restart, with max F crashes at a time	$\mathcal{P}_F^{recovery} = \bigotimes_{j=1}^F \mathcal{P}_1^{recovery}$	$HOProd(\{T \subseteq \Pi \mid T \geq n-F\})$	Thm 4
At most 1 crash, after round r	$\mathcal{P}_{1,\geq r}^{crash} = \bigcup_{i=r}^{\infty} \mathcal{P}_{1,i}^{crash}$	$\subseteq HOProd(\{T \subseteq \Pi \mid T \geq n-1\})$	Thm 10
At most F crashes, after round r	$\mathcal{P}_{F,\geq r}^{crash} = \bigotimes_{i=r}^{\infty} \mathcal{P}_{F,i}^{crash}$	$\subseteq HOProd(\{T \subseteq \Pi \mid T \geq n-F\})$	Thm 10
At most F crashes with no more than one per round	$\mathcal{P}_F^{crash\neq} = \bigcup_{i_1 \neq i_2 \neq \dots \neq i_F} \bigotimes_{j=1}^F \mathcal{P}_{1,i_j}^{crash}$	$\subseteq HOProd(\{T \subseteq \Pi \mid T \geq n-F\})$	Thm 10

Table 1. A list of delivered predicate built using our operations, and their corresponding heard-of predicate. The *HOProduct* operator is defined in Definition 15.

2.4 Examples

We now show the variety of models that can be constructed from basic building blocks. Our basic blocks are the model $PDel^{total}$ with only the collection c_{total} where every message is delivered, and the model $PDel_{1,r}^{crash}$ with at most one crash that can happen at round r .

Definition 11 (At most 1 crash at round r). $\mathcal{P}_{1,r}^{crash} \triangleq$

$$\left\{ c \in (\mathbb{N}^* \times \Pi) \mapsto \mathcal{P}(\Pi) \mid \exists \Sigma \subseteq \Pi : \begin{array}{l} |\Sigma| \geq n-1 \\ \wedge \forall j \in \Pi \left(\begin{array}{l} \forall r' \in [1, r[: c(r', j) = \Pi \\ \wedge c(r, j) \supseteq \Sigma \\ \wedge \forall r' \geq r : c(r', j) = \Sigma \end{array} \right) \end{array} \right\}.$$

From this family of predicates, various predicates can be built (Table 1).

3 Oblivious Strategies

The simplest strategies use only information from the messages of the current round. These strategies that do not remember messages from previous rounds, do not use messages in advance from future rounds, and do not use the round number itself.

We call these strategies oblivious. They are simple, the Heard-Of predicates they implement are relatively easy to compute, and they require little computing power and memory to implement. Moreover, many examples above are dominated by such a strategy. Of course, there is a price to pay: oblivious strategies tend to be coarser than general ones.

3.1 Minimal Oblivious Strategy

An oblivious strategy is defined by the different subsets of Π from which it has to receive a message before allowing a change of round.

Definition 12 (Oblivious Strategy). *Let f be a strategy and, $\forall q \in Q$, let $obliv(q) = \{k \in \Pi \mid \langle i, k \rangle \in q.mes\}$. f is an **oblivious strategy** $\triangleq \forall q, q' \in Q : obliv(q) = obliv(q') \implies (q \in f \iff q' \in f)$. For f an oblivious strategy, let $Nests_f \triangleq \{obliv(q) \mid q \in f\}$. It uniquely defines f .*

We will focus on a specific strategy, that dominates the oblivious strategies for a predicate. This follows from the fact that it waits less than any other valid oblivious strategy for this predicate.

Definition 13 (Minimal Oblivious Strategy). *Let $PDel$ be a delivered predicate. The **minimal oblivious strategy** for $PDel$ is $f_{min} \triangleq \{q \mid \exists c \in PDel, \exists p \in \Pi, \exists r > 0 : obliv(q) = c(r, p)\}$.*

Lemma 1 (Domination of Minimal Oblivious Strategy). *Let $PDel$ be a $PDel$ and f_{min} be its minimal oblivious strategy. Then f_{min} is a dominating oblivious strategy for $PDel$.*

Proof (Proof idea). f_{min} is valid, because for every possible set of received messages in a collection of $PDel$, it accepts the corresponding oblivious state by definition of minimal oblivious strategy. It is dominating among oblivious strategies because every other valid oblivious strategy must allow the change of round when f_{min} does it: it contains f_{min} . If an oblivious strategy does not contain f_{min} , then there is a collection of $PDel$ in which at a given round, a certain process might receive exactly the messages for the oblivious state accepted by f_{min} and not by f . This entails that f is not valid.

3.2 Operations Maintain Minimal Oblivious Strategy

As teased above, minimal oblivious strategies behave nicely under the proposed operations. That is, they give minimal oblivious strategies of resulting delivered

predicates. One specificity of minimal oblivious strategies is that there is no need for the succession operation on strategies, nor for the repetition. An oblivious strategy has no knowledge about anything but the messages of the current round, and not even its round number, so it is impossible to distinguish a union from a succession, or a repetition from the initial predicate itself.

Theorem 1 (Minimal Oblivious Strategy for Union and Succession). *Let $PDel_1, PDel_2$ be two delivered predicates, f_1 and f_2 the minimal oblivious strategies for, respectively, $PDel_1$ and $PDel_2$. Then $f_1 \cup f_2$ is the minimal oblivious strategy for $PDel_1 \cup PDel_2$ and $PDel_1 \rightsquigarrow PDel_2$.*

Proof (Proof idea). Structurally, all proofs in this section consist in showing equality between the strategies and the minimal oblivious strategy for the delivered predicate. In a union, the strategy can never know in which $PDel$ it is. Similarly it can never know if the succession happened already or not in the succession. Thus, the minimal oblivious strategy must wait for any set of messages present in any collection, that is any set of messages in a collection of $PDel_1$ or $PDel_2$. Given that f_1 and f_2 are the minimal oblivious strategies of $PDel_1$ and $PDel_2$, they accept exactly the states with one of these sets of current messages. And thus $f_1 \cup f_2$ is the minimal oblivious strategy for $PDel_1 \cup PDel_2$ and $PDel_1 \rightsquigarrow PDel_2$.

Theorem 2 (Minimal Oblivious Strategy for Repetition). *Let $PDel$ be a delivered predicate, and f be its minimal oblivious strategy. Then f is the minimal oblivious strategy for $PDel^\omega$.*

Proof (Proof idea). The intuition is the same as for union and succession. Since repetition involves only one $PDel$, the sets of received messages do not change and f is the minimal oblivious strategy.

For combination, an additional hypothesis is needed: in addition to symmetry over processes, we need symmetry over rounds, because oblivious strategies cannot know the round.

Definition 14 (Totally Symmetric $PDel$). *Let $PDel$ be a delivered predicate. $PDel$ is **totally symmetric** $\triangleq \forall c \in PDel, \forall r > 0, \forall p \in \Pi, \forall r' > 0, \forall q \in \Pi, \exists c' \in PDel : c(r, p) = c'(r', q)$*

Combination is different because combining collections is done round by round. As oblivious strategies do not depend on the round, the combination of oblivious strategies creates the same combination of received messages for each round. We thus need these combinations to be independent of the round – to be possible at each round – to reconcile those two elements.

Theorem 3 (Minimal Oblivious Strategy for Combination). *Let $PDel_1, PDel_2$ be two totally symmetric delivered predicates, f_1 and f_2 the minimal oblivious strategies for, respectively, $PDel_1$ and $PDel_2$. Then $f_1 \otimes f_2$ is the minimal oblivious strategy for $PDel_1 \otimes PDel_2$.*

Proof (Proof idea). The oblivious states of $PDel_1 \otimes PDel_2$ are the combination of an oblivious state of $PDel_1$ and of one of $PDel_2$ at the same round, for the same process. Thanks to total symmetry, this translates into the intersection of any oblivious state of $PDel_1$ with any oblivious state of $PDel_2$. Since f_1 and f_2 are the minimal oblivious strategy, they both accept exactly the oblivious states of $PDel_1$ and $PDel_2$ respectively. Thus, $f_1 \otimes f_2$ accept all combinations of oblivious states of $PDel_1$ and $PDel_2$, and thus is the minimal oblivious strategy of $PDel_1 \otimes PDel_2$.

3.3 Computing Heard-Of Predicates

The computation of the heard-of predicate generated by an oblivious strategy is easy thanks to a characteristic of this HO: it is a product of sets of possible messages.

Definition 15 (Heard-Of Product). *Let $S \subseteq \mathcal{P}(H)$. The **heard-of product generated by S** , $HOProd(S) \triangleq \{h \mid \forall p \in H, \forall r > 0 : h(r, p) \in S\}$.*

Lemma 2 (Heard-Of Predicate of an Oblivious Strategy). *Let $PDel$ be a delivered predicate containing c_{tot} and let f be a valid oblivious strategy for $PDel$. Then $PHO_f(PDel) = HOProd(Nexts_f)$.*

Proof. Proved in [12, Theorem 20, Section 4.1].

Thanks to this characterization, the heard-of predicate generated by the minimal strategies for the operations is computed in terms of the heard-of predicate generated by the original minimal strategies.

Theorem 4 (Heard-Of Predicate of Minimal Oblivious Strategies). *Let $PDel, PDel_1, PDel_2$ be delivered predicates containing c_{tot} . Let f, f_1, f_2 be their respective minimal oblivious strategies. Then:*

- $PHO_{f_1 \cup f_2}(PDel_1 \cup PDel_2) = PHO_{f_1 \cup f_2}(PDel_1 \rightsquigarrow PDel_2)$
 $= HOProd(Nexts_{f_1} \cup Nexts_{f_2})$.
- If $PDel_1$ or $PDel_2$ are totally symmetric, $PHO_{f_1 \otimes f_2}(PDel_1 \otimes PDel_2) = HOProd(\{n_1 \cap n_2 \mid n_1 \in Nexts_{f_1} \wedge n_2 \in Nexts_{f_2}\})$.
- $PHO_f(PDel^\omega) = PHO_f(PDel)$.

Proof (Proof idea). We apply Lemma 2. The containment of c_{tot} was shown in the proof of Theorem 5. As for the equality of the oblivious states, it follows from the intuition in the proofs of the minimal oblivious strategy in the previous section.

3.4 Domination by an Oblivious Strategy

From the previous sections, we can compute the Heard-Of predicate of the dominating oblivious strategies for our examples. We first need to give the minimal oblivious strategy for our building blocks $PDel_1^{crash}$ and $PDel^{total}$.

Definition 16 (Waiting for $n - F$ messages). *The strategy to wait for $n - F$ messages is: $f^{n,F} \triangleq \{q \in Q \mid |\text{obliv}(q)| \geq n - F\}$*

For all $F < n$, $f^{n,F}$ is the minimal oblivious strategy for $PDel_F^{crash}$ (shown by Shimi et al. [12, Thm. 17]). For $PDel^{total}$, since every process receives all the messages all the time, the strategy waits for all the messages ($f^{n,0}$).

Using these strategies, we deduce the heard-of predicates of dominating oblivious strategies for our examples.

- For $PDel_1^{recover} \triangleq PDel_1^{crash} \rightsquigarrow PDel^{total}$, the minimal oblivious strategy $f_1^{recover} = f^{n,1} \cup f^{n,0} = f^{n,1}$. This entails that $PHO_{f_1^{recover}} = HOProd(\{T \subseteq \Pi \mid |T| \geq n - 1\})$.
- For $PDel_1^{canrecover} \triangleq PDel_1^{recover} \cup PDel_1^{crash}$, the minimal oblivious strategy $f_1^{canrecover} = f_1^{recover} \cup f^{n,1} = f^{n,1}$. This entails that $PHO_{f_1^{canrecover}} = HOProd(\{T \subseteq \Pi \mid |T| \geq n - 1\})$.
- For $PDel_1^{crash} \otimes PDel_1^{canrecover}$ the minimal oblivious strategy $f = f^{n,1} \otimes f_1^{canrecover} = f^{n,1} \otimes f^{n,1} = f^{n,2}$. This entails that $PHO_f = HOProd(\{T \subseteq \Pi \mid |T| \geq n - 2\})$.

The computed predicate is the predicate of the dominating *oblivious* strategy. But the dominating strategy might not be oblivious, and this predicate might be too weak. The following result shows that $PDel_1^{crash}$ and $PDel^{total}$ satisfy conditions that imply their domination by an oblivious strategy. Since these conditions are invariant by our operations, all $PDel$ constructed with these building blocks are dominated by an oblivious strategy.

Theorem 5 (Domination by Oblivious for Operations).

Let $PDel_1, PDel_2$ be delivered predicates that satisfy:

- (**Total collection**) *They contains the total collection c_{tot} ,*
- (**Symmetry up to a round**) $\forall c$ a collection in the predicate, $\forall p \in \Pi, \forall r > 0, \forall r' > 0, \exists c'$ a collection in the predicate: $c'[1, r' - 1] = c_{tot}[1, r' - 1] \wedge \forall q \in \Pi : c'(r', q) = c(r, p)$

Then $PDel_1 \cup PDel_2, PDel_1 \otimes PDel_2, PDel_1 \rightsquigarrow PDel_2, PDel^\omega$ satisfy the same two conditions and are dominated by oblivious strategies.

Both \mathcal{P}_1^{crash} from Table 1 and $\mathcal{P}^{total} = \{c_{tot}\}$ satisfy this condition. So do all the first 8 examples from Table 1, since they are built from these two.

4 Conservative Strategies

We now broaden our class of considered strategies, by allowing them to consider past and present rounds, as well as the round number itself. This is a generalization of oblivious strategies, that tradeoff simplicity for expressivity, while retaining a nice structure. Even better, we show that both our building blocks and all the predicates built from them are dominated by such a strategy. For the examples then, no expressivity is lost.

4.1 Minimal Conservative Strategy

Definition 17 (Conservative Strategy). Let f be a strategy, and $\forall q \in Q$, let $\text{cons}(q) \triangleq \langle q.\text{round}, \{\langle r, k \rangle \in q.\text{mes} \mid r \leq q.\text{round}\} \rangle$. Then f is a **conservative strategy** $\triangleq \forall q, q' \in Q : \text{cons}(q) = \text{cons}(q') \implies (q \in f \iff q' \in f)$.

We write $\text{Nexts}_f^R \triangleq \{\text{cons}(q) \mid q \in f\}$ for the set of conservative states in f . This uniquely defines f .

In analogy with the case of oblivious strategies, we can define a minimal conservative strategy of $PDel$, and it is a strategy dominating all conservative strategies for this delivered predicate.

Definition 18 (Minimal Conservative Strategy). Let $PDel$ be a delivered predicate. Then the **minimal conservative strategy** for $PDel$ is $f_{min} \triangleq$ the conservative strategy such that $f = \{q \in Q \mid \exists c \in PDel, \exists p \in \Pi, \forall r \leq q.\text{round} : q(r) = c(r, p)\}$.

Lemma 3 (Domination of Minimal Conservative Strategy). Let $PDel$ be a delivered predicate and f_{min} be its minimal conservative strategy. Then f_{min} dominates the conservative strategies for $PDel$.

Proof (Proof idea). Analogous to the case of minimal oblivious strategies: it is valid because it allow to change round for each possible conservative state (the round and the messages received for this round and before) of collections in $PDel$. And since every other valid conservative strategy f must accept these states (or they would block forever in some execution of a collection of $PDel$), we have that f contains f_{min} and thus that f_{min} dominates f .

4.2 Operations Maintain Minimal Conservative Strategies

Like oblivious strategies, minimal conservative strategies give minimal conservative strategies of resulting delivered predicates.

Theorem 6 (Minimal Conservative Strategy for Union).

Let $PDel_1, PDel_2$ be two delivered predicates, f_1 and f_2 the minimal conservative strategies for, respectively, $PDel_1$ and $PDel_2$. Then $f_1 \cup f_2$ is the minimal conservative strategy for $PDel_1 \cup PDel_2$.

Proof (Proof idea). Each prefix of a collection in $PDel_1 \cup PDel_2$ comes from either $PDel_1$ or $PDel_2$, and thus is accepted by f_1 or f_2 . And any state accepted by $f_1 \cup f_2$ corresponds to some prefix of $PDel_1$ or $PDel_2$.

For the other three operations, slightly more structure is needed on the predicates. More precisely, they have to be independent of the processes. Any prefix of a process p in a collection of the predicate is also the prefix of any other process q in a possibly different collection of the same $PDel$. Hence, the behaviors (fault, crashes, loss) are not targeting specific processes. This restriction fits the intuition behind many common fault models.

Definition 19 (Symmetric PDel). Let $PDel$ be a delivered predicate. $PDel$ is **symmetric** $\triangleq \forall c \in PDel, \forall p \in \Pi, \forall r > 0, \forall q \in \Pi, \exists c' \in PDel, \forall r' \leq r : c'(r', q) = c(r', p)$

Theorem 7 (Minimal Conservative Strategy for Combination).

Let $PDel_1, PDel_2$ be two symmetric delivered predicates, f_1 and f_2 the minimal conservative strategies for, respectively, $PDel_1$ and $PDel_2$. Then $f_1 \otimes f_2$ is the minimal conservative strategy for $PDel_1 \otimes PDel_2$.

Proof (Proof idea). Since f_1 and f_2 are the minimal conservative strategies of $PDel_1$ and $PDel_2$, $Nexts^R f_1$ is the set of the conservative states of prefixes of $PDel_1$ and $Nexts^R_{f_2}$ is the set of the conservative states of prefixes of $PDel_2$. Also, the states accepted by $f_1 \otimes f_2$ are the combination of the states accepted by f_1 and the states accepted by f_2 . And the prefixes of $PDel_1 \otimes PDel_2$ are the prefixes of $PDel_1$ combined with the prefixes of $PDel_2$ **for the same process**. Thanks to symmetry, we can take a prefix of $PDel_2$ and any process, and find a collection such that the process has that prefix. Therefore the combined prefixes for the same process are the same as the combined prefixes of $PDel_1$ and $PDel_2$. Thus, $Nexts^R_{f_1 \otimes f_2}$ is the set of conservative states of prefixes of $PDel_1 \otimes PDel_2$, and $f_1 \otimes f_2$ is its minimal conservative strategy.

Theorem 8 (Minimal Conservative Strategy for Succession).

Let $PDel_1, PDel_2$ be two symmetric delivered predicates, f_1 and f_2 the minimal conservative strategies for, respectively, $PDel_1$ and $PDel_2$. Then $f_1 \rightsquigarrow f_2$ is the minimal conservative strategy for $PDel_1 \rightsquigarrow PDel_2$.

Proof (Proof idea). Since f_1 and f_2 are the minimal conservative strategies of $PDel_1$ and $PDel_2$, $Nexts^R f_1$ is the set of the conservative states of prefixes of $PDel_1$ and $Nexts^R_{f_2}$ is the set of the conservative states of prefixes of $PDel_2$. Also, the states accepted by $f_1 \rightsquigarrow f_2$ are the succession of the states accepted by f_1 and the states accepted by f_2 . And the prefixes of $PDel_1 \rightsquigarrow PDel_2$ are the successions of prefixes of $PDel_1$ and prefixes of $PDel_2$ **for the same process**. But thanks to symmetry, we can take a prefix of $PDel_2$ and any process, and find a collection such that the process has that prefix.

Therefore the succession of prefixes for the same process are the same as the succession of prefixes of $PDel_1$ and $PDel_2$. Thus, $Nexts^R_{f_1 \rightsquigarrow f_2}$ is the set of conservative states of prefixes of $PDel_1 \rightsquigarrow PDel_2$, and is therefore its minimal conservative strategy.

Theorem 9 (Minimal Conservative Strategy for Repetition).

Let $PDel$ be a symmetric delivered predicate, and f be its minimal conservative strategy. Then f^ω is the minimal conservative strategy for $PDel^\omega$.

Proof (Proof idea). The idea is the same as in the succession.

4.3 Computing Heard-Of Predicates

Here we split from the analogy with oblivious strategies: the heard-of predicate of conservative strategies is hard to compute, as it depends in intricate ways on the delivered predicate itself.

Yet it is still possible to compute interesting information on this HO: upper bounds. These are overapproximations of the actual HO, but they can serve for formal verification.

Theorem 10 (Upper Bounds on HO of Minimal Conservative Strategies). *Let $PDel, PDel_1, PDel_2$ be delivered predicates containing c_{tot} .*

Let $f^{cons}, f_1^{cons}, f_2^{cons}$ be their respective minimal conservative strategies, and $f^{obliv}, f_1^{obliv}, f_2^{obliv}$ be their respective minimal oblivious strategies. Then:

- $PHO_{f_1^{cons} \cup f_2^{cons}}(PDel_1 \cup PDel_2) \subseteq HOProd(Nexts_{f_1^{obliv}} \cup Nexts_{f_2^{obliv}})$.
- $PHO_{f_1^{cons} \rightsquigarrow f_2^{cons}}(PDel_1 \rightsquigarrow PDel_2) \subseteq HOProd(Nexts_{f_1^{obliv}} \cup Nexts_{f_2^{obliv}})$.
- $PHO_{f_1^{cons} \otimes f_2^{cons}}(PDel_1 \otimes PDel_2) \subseteq HOProd(\{n_1 \cap n_2 \mid n_1 \in Nexts_{f_1^{obliv}} \wedge n_2 \in Nexts_{f_2^{obliv}}\})$.
- $PHO_{(f^{cons})^\omega}(PDel^\omega) \subseteq HOProd(Nexts_{f^{obliv}})$.

Proof (Proof idea). These bounds follow from the fact that an oblivious strategy, is a conservative strategy, and thus the minimal conservative strategy dominates the minimal oblivious strategy.

5 Conclusion

To summarize, we propose operations on delivered predicates that allow the construction of complex predicates from simpler ones. The corresponding operations on strategies behave nicely regarding dominating strategies, for the conservative and oblivious strategies. This entails bounds and characterizations of the dominating heard-of predicate for the constructions.

What needs to be done next comes in two kinds: first, the logical continuation is to look for constraints on delivered predicates for which we can compute the dominating heard-of predicate of conservative strategies. More ambitiously, we will study strategies looking in the future, i.e. strategies that can take into account messages from processes that have already reached a strictly higher round than the recipient. These strategies are useful for inherently asymmetric delivered predicates. For example, message loss is asymmetric, in the sense that we cannot force processes to receive the same set of messages.

References

1. Martin Biely, Peter Robinson, Manfred Schmid, Ulrich Schwarz, and Kyrill Winkler. Gracefully degrading consensus and k-set agreement in directed dynamic networks. *Theoretical Computer Science*, 726:41–77, 2018. doi:10.1016/j.tcs.2018.02.019.
2. Bernadette Charron-Bost, Henri Debrat, and Stephan Merz. Formal verification of consensus algorithms tolerating malicious faults. In *Stabilization, Safety, and Security of Distributed Systems*, pages 120–134, 2011. doi:10.1007/978-3-642-24550-3_11.

3. Bernadette Charron-Bost, Matthias Függer, and Thomas Nowak. Approximate consensus in highly dynamic networks: The role of averaging algorithms. In *Automata, Languages, and Programming*, pages 528–539, 2015. doi:10.1007/978-3-662-47666-6_42.
4. Bernadette Charron-Bost and André Schiper. The heard-of model: computing in distributed systems with benign faults. *Distributed Computing*, 22(1):49–71, April 2009. doi:10.1007/s00446-009-0084-6.
5. Étienne Coulouma, Emmanuel Godard, and Joseph Peters. A characterization of oblivious message adversaries for which consensus is solvable. *Theoretical Computer Science*, 584:80–90, 2015. doi:10.1016/j.tcs.2015.01.024.
6. Cezara Drăgoi, Thomas A. Henzinger, and Damien Zufferey. Psync: A partially synchronous language for fault-tolerant distributed algorithms. *SIGPLAN Not.*, 51(1):400–415, January 2016. doi:10.1145/2914770.2837650.
7. Tzilla Elrad and Nissim Francez. Decomposition of distributed programs into communication-closed layers. *Science of Computer Programming*, 2(3):155–173, 1982. doi:10.1016/0167-6423(83)90013-8.
8. Eli Gafni. Round-by-round fault detectors (extended abstract): Unifying synchrony and asynchrony. In *17th ACM Symposium on Principles of Distributed Computing*, PODC '98, pages 143–152, 1998. doi:10.1145/277697.277724.
9. Ognjen Marić, Christoph Sprenger, and David Basin. Cutoff bounds for consensus algorithms. In *Computer Aided Verification*, pages 217–237, 2017. doi:10.1007/978-3-319-63390-9_12.
10. Thomas Nowak, Ulrich Schmid, and Kyrill Winkler. Topological characterization of consensus under general message adversaries. In *2019 ACM Symposium on Principles of Distributed Computing*, PODC '19, 2019. doi:10.1145/3293611.3331624.
11. Nicola Santoro and Peter Widmayer. Time is not a healer. In *6th Symposium on Theoretical Aspects of Computer Science STACS 89*, pages 304–313, 1989. doi:10.1007/BFb0028994.
12. Adam Shimi, Aurélie Hurault, and Philippe Quéinnec. Characterizing asynchronous message-passing models through rounds. In *22nd Int'l Conf. on Principles of Distributed Systems (OPODIS 2018)*, pages 18:1–18:17, 2018. doi:10.4230/LIPIcs.OPODIS.2018.18.

A Tools

A.1 Timing Functions

A timing function of an execution captures the round at which a message is delivered: for a message sent in round r' by k to j , $time(r', k, j)$ is the round at which this message is delivered to j . Note that $time(r', k, j) = 0$ if and only if no message sent from k to j at round r' is delivered in this execution.

Definition 20 (Timing Function). A *timing function* is a function $\mathbb{N}^* \times \Pi \times \Pi \mapsto \mathbb{N}$.

For t an execution, the *timing function of t* , $time_t \triangleq$ the timing function such that $\forall r > 0, \forall r' > 0, \forall k, j \in \Pi : time_t(r', k, j) = r \iff (\exists i \geq 0 : t[i] = deliver(r', k, j) \wedge q_j^t[i].round = r)$.

The standard execution reorders deliveries and changes of round such that all the deliveries for a given round happen before the changes of round for all processes.

Definition 21 (Standard Execution of a timing function). Let $time$ be a timing function and ord be any function taking a set and returning an ordered sequence of its elements. The specific ordering doesn't matter.

The *standard execution with timing $time$* is $st_{time} \triangleq \prod_{r \in \mathbb{N}^*} dels_r.nexts$, where $dels_r = ord(\{deliver(r', k, j) \mid r' > 0 \wedge k, j \in \Pi \wedge time(r', k, j) = r\})$ and $nexts = ord(\{next_p \mid p \in \Pi\})$.

Lemma 4 (Correctness of Standard Execution with Timing). Let $time$ be a timing function. Then $(\forall r > 0, \forall k, j \in \Pi : time(r, k, j) = 0 \vee time(r, k, j) \geq r) \implies st_{time}$ is an execution.

Proof. – **(Delivered after sending)** Let $r > 0$ and $k, j \in \Pi$. If $time(r, k, j) = 0$, then the message is never delivered, and we don't have to consider it. If not, then by hypothesis $time(r, k, j) \geq r$. This means $\exists i \geq r : deliver(r, k, j) \in dels_i$.

By construction of the standard execution, there are $i - 1$ occurrences of the sequence $nexts$ before the sequence $dels_i$. This means there are $i - 1 \geq r - 1$ $next_k$ before, which allows us to conclude.

- **(Delivered only once)** Let $r > 0$ and $k, j \in \Pi$. If $\exists i \geq 0 : st_{time}[i] = deliver(r, k, j)$, then it is in $dels_{time(r, k, j)}$. We conclude that there is only one delivery of this message.
- **(Once stopped, forever stopped)** The standard execution does not contain any *stop*.

Lemma 5 (Heard-Of Collection of Timing Function). Let t be a valid execution, and $time$ be its timing function. Then $\forall r > 0, \forall p \in \Pi : h_t(r, p) = \{q \in \Pi \mid time(r, q, p) \in [1, r]\}$.

Proof. Let $i \geq 0$ such that $\exists p \in \Pi : t[i] = next_p$. Let $r = q_p^t[i].round$. We show both side of $h_t(r, p) = \{q \in \Pi \mid time(r, q, p) \in [1, r]\}$.

- Let $q \in h_t(r, p)$. Then it is delivered in a round $\leq r$, and thus $time(r, q, p) \in [1, r]$.
- Let $q \in \Pi$ such that $time(r, q, p) \in [1, r]$. Then by definition of $time$, the message sent by q at round r is delivered to p in t at most at round r . Thus, it is in the messages from the current round when going to round $r + 1$, and $q \in h_t(r, p)$.

B Proofs for Oblivious Strategies

B.1 Minimal Oblivious Strategies

We use a necessary and sufficient condition for an oblivious strategy to be valid in the rest of the proofs.

Lemma 6 (Necessary and Sufficient Condition for Validity of a Oblivious Strategy). *Let $PDel$ be a delivered predicate and f be an oblivious strategy. Then f is valid for $PDel \iff f \supseteq \{q \mid \exists c \in PDel, \exists p \in \Pi, \exists r > 0 : obliv(q) = c(r, p)\}$.*

Proof. From the version in OPODIS 2018, f has to satisfy $\forall c \in PDel, \forall r > 0, \forall p \in \Pi : c(r, p) \in Nexts_f$.

We show the equivalence of this condition with our own, which allow us to conclude by transitivity of equivalence.

- (\implies) We assume our condition holds and prove the one from OPODIS 2018. Let $c \in PDel, r > 0$ and $p \in \Pi$: we want to show that $c(r, p) \in Nexts_f$. That is to say, that all states whose present corresponds to this oblivious state are accepted by f . Let q such that $obliv(q) = c(r, p)$. We have the collection c , the round r and the process p to apply our condition, and thus $q \in f$. Hence, $c(r, p) \in Nexts_f$.
- (\impliedby) We assume the condition from OPODIS 2018 holds and we prove ours. Let q such that $\exists c \in PDel, \exists p \in \Pi, \exists r \leq q.round : obliv(q) = c(r, p)$. By hypothesis, we have $c(r, p) \in Nexts_f$. We conclude that $q \in f$.

Lemma ((1 Domination of Minimal Oblivious Strategy)). *Let $PDel$ be a $PDel$ and f_{min} be its minimal oblivious strategy. Then f_{min} is a dominating oblivious strategy for $PDel$.*

Proof. First, f_{min} is valid for $PDel$ by application of Lemma 6. Next, we take another oblivious strategy f , which is valid for $PDel$. Lemma 6 now gives us that $f_{min} \subseteq f$. Hence, when f_{min} allow a change of round, so does f . This entails that all executions of f_{min} for $PDel$ are also executions of f for $PDel$, and thus that heard-of predicate generated by f_{min} is contained in the one generated by f .

B.2 Operations Maintain Minimal Oblivious Strategies

Theorem ((1) Minimal Oblivious Strategy for Union and Succession). *Let $PDel_1, PDel_2$ be two delivered predicates, f_1 the minimal oblivious strategy for $PDel_1$, and f_2 the minimal oblivious strategy for $PDel_2$. Then $f_1 \cup f_2$ is the minimal oblivious strategy for $PDel_1 \cup PDel_2$ and $PDel_1 \rightsquigarrow PDel_2$.*

Proof. We first show that the minimal oblivious strategies of $PDel_1 \cup PDel_2$ and $PDel_1 \rightsquigarrow PDel_2$ are equal. Hence, we prove $\{q \mid \exists c \in PDel_1 \cup PDel_2, \exists p \in \Pi, \exists r > 0 : \text{obliv}(q) = c(r, p)\} = \{q \mid \exists c \in PDel_1 \rightsquigarrow PDel_2, \exists p \in \Pi, \exists r > 0 : \text{obliv}(q) = c(r, p)\}$.

- (\subseteq) Let q such that $\exists c \in PDel_1 \cup PDel_2, \exists p \in \Pi, \exists r > 0 : \text{obliv}(q) = c(r, p)$.
 - If $c \in PDel_1$, then we take $c_2 \in PDel_2$ $c' = c[1, r].c_2$. Since $c' \in c \rightsquigarrow c_2$, we have $c' \in PDel_1 \rightsquigarrow PDel_2$. And by definition of c' , $c'(r, p) = c(r, p)$. We thus have c', p and r showing that q is in the set on the right.
 - If $c \in PDel_2$, then $c \in PDel_1 \rightsquigarrow PDel_2$. We thus have c, p and r showing that q is in the set on the right.
- (\supseteq) Let q such that $\exists c \in PDel_1 \rightsquigarrow PDel_2, \exists p \in \Pi, \exists r > 0 : \text{obliv}(q) = c(r, p)$.
 - If $c \in PDel_2$, then $c \in PDel_1 \cup PDel_2$. We thus have c, p and r showing that q is in the set on the left.
 - If $c \notin PDel_2$, there exist $c_1 \in PDel_1, c_2 \in PDel_2$ and $r' > 0$ such that $c = c_1[1, r'].c_2$.
 - * If $r \leq r'$, then by definition of c , we have $c(r, p) = c_1(r, p)$. We thus have c_1, p and r showing that q is in the set on the left.
 - * If $r > r'$, then $c(r, p) = c_2(r - r', p)$. We thus have c_2, p and $(r - r')$ showing that q is in the set on the left.

We show that $f_1 \cup f_2 = \{q \mid \exists c \in PDel_1 \cup PDel_2, \exists p \in \Pi, \exists r > 0 : \text{obliv}(q) = c(r, p)\}$, which allows us to conclude by Definition 13.

- Let $q \in f_1 \cup f_2$. We fix $q \in f_1$ (the case $q \in f_2$ is completely symmetric). Then because f_1 is the minimal oblivious strategy of $PDel_1$, by application of Lemma 6, $\exists c_1 \in PDel_1, \exists p \in \Pi, \exists r > 0$ such that $c_1(r, p) = \text{obliv}(q)$. $c_1 \in PDel_1 \subseteq PDel_1 \cup PDel_2$. We thus have c_1, p and r showing that q is in the minimal oblivious strategy for $PDel_1 \cup PDel_2$.
- Let q such that $\exists c \in PDel_1 \cup PDel_2, \exists p \in \Pi, \exists r > 0 : c(r, p) = \text{obliv}(q)$. By definition of union, c must be in $PDel_1$ or in $c \in PDel_2$; we fix $c \in PDel_1$ (the case $PDel_2$ is symmetric). Then Definition 13 gives us that q is in the minimal oblivious strategy of $PDel_1$, that is f_1 . We conclude that $q \in f_1 \cup f_2$.

Theorem ((2) Minimal Oblivious Strategy for Repetition). *Let $PDel$ be a delivered predicate, and f be its minimal oblivious strategy. Then f is the minimal oblivious strategy for $PDel^\omega$.*

Proof. We show that $f = \{q \mid \exists c \in PDel^\omega, \exists p \in \Pi, \exists r > 0 : \text{obliv}(q) = c(r, p)\}$, which allows us to conclude by Definition 13.

- (\sqsubseteq) Let $q \in f$. By minimality of f for $PDel$, $\exists c \in PDel, \exists p \in \Pi, \exists r > 0 : \text{obliv}(q) = c(r, p)$.
We take $c' \in PDel^\omega$ such that $c_1 = c$ and $r_2 = r$; the other c_i and r_i don't matter for the proof. By definition of repetition, we get $c'(r, p) = c(r, p) = \text{obliv}(q)$.
We have c', p and r showing that q is in the minimal oblivious strategy of $PDel^\omega$.
- (\supseteq) Let q such that $\exists c \in PDel^\omega, \exists p \in \Pi, \exists r > 0 : \text{obliv}(q) = c(r, p)$. By definition of repetition, there are $c_i \in PDel$ and $0 < r_i < r_{i+1}$ such that $r \in [r_i + 1, r_{i+1}]$ and $c(r, p) = c_i(r - r_i, p)$.
We have found c_i, p and $(r - r_i)$ showing that q is in the minimal oblivious strategy for $PDel$. And since f is the minimal oblivious strategy for $PDel$, we get $q \in f$.

Theorem ((3) Minimal Oblivious Strategy for Combination). *Let $PDel_1, PDel_2$ be two totally symmetric delivered predicate, f_1 the minimal oblivious strategy for $PDel_1$, and f_2 the minimal oblivious strategy for $PDel_2$. Then $f_1 \otimes f_2$ is the minimal oblivious strategy for $PDel_1 \otimes PDel_2$.*

Proof. We show that $f_1 \otimes f_2 = \{q \mid \exists c \in PDel_1 \otimes PDel_2, \exists p \in \Pi, \exists r > 0 : \text{obliv}(q) = c(r, p)\}$, which allows us to apply Lemma 13.

- Let $q \in f_1 \otimes f_1$. Then $\exists q_1 \in f_1, \exists q_2 \in f_2$ such that $q = q_1 \otimes q_2$. This also means that $q_1.\text{round} = q_2.\text{round} = q.\text{round}$.
By minimality of f_1 and f_2 , $\exists c_1 \in PDel_1, \exists p_1 \in \Pi, \exists r_1 > 0 : c_1(r_1, p_1) = \text{obliv}(q_1)$ and $\exists c_2 \in PDel_2, \exists p_2 \in \Pi, \exists r_2 > 0 : c_2(r_2, p_2) = \text{obliv}(q_2)$.
Moreover, total symmetry of $PDel_2$ ensures that $\exists c'_2 \in PDel_2 : c'_2(r_1, p_1) = c_2(r_2, p_2)$.
We take $c = c_1 \otimes c'_2$. $\text{obliv}(q) = \text{obliv}(q_1) \cap \text{obliv}(q_2) = c_1(r_1, p_1) \cap c_2(r_2, p_2) = c_1(r_1, p_1) \cap c'_2(r_1, p_1) = c(r_1, p_1)$.
We have c, p_1 and r_1 showing that q is in the minimal oblivious strategy for $PDel_1 \otimes PDel_2$.
- Let q such that $\exists c \in PDel_1 \otimes PDel_2, \exists p \in \Pi, \exists r > 0 : c(r, p) = \text{obliv}(q)$. By definition of Combination, $\exists c_1 \in PDel_1, \exists c_2 \in PDel_2 : c = c_1 \otimes c_2$.
We take q_1 such that $q_1.\text{round} = r, \text{obliv}(q_1) = c_1(r, p)$ and $\forall r' \neq r : q_1(r') = q(r')$; we also take q_2 such that $q_2.\text{round} = r, \text{obliv}(q_2) = c_2(r, p)$ and $\forall r' \neq r : q_2(r') = q(r')$.
Then $q = q_1 \otimes q_2$. And since f_1 and f_2 are the minimal oblivious strategies of $PDel_1$ and $PDel_2$ respectively, we have $q_1 \in f_1$ and $q_2 \in f_2$.
We conclude that $q \in f_1 \otimes f_2$.

B.3 Computing Heard-Of Predicates

Theorem ((4) Heard-Of Predicate of Minimal Oblivious Strategies). *Let $PDel, PDel_1, PDel_2$ be delivered predicates containing c_{tot} . Let HO, HO_1, HO_2 be their respective HO , and let f, f_1, f_2 be their respective minimal oblivious strategies. Then:*

- The HO generated by $f_1 \cup f_2$ on $PDel_1 \cup PDel_2$, and on $PDel_1 \rightsquigarrow PDel_2$ is the HO product generated by $Nexts_{f_1} \cup Nexts_{f_2}$.
- The HO generated by $f_1 \otimes f_2$ on $PDel_1 \otimes PDel_2$, if either $PDel_1$ or $PDel_2$ is totally symmetric, is the HO product generated by $\{n_1 \cap n_2 \mid n_1 \in Nexts_{f_1} \wedge n_2 \in Nexts_{f_2}\}$.
- The HO generated by f on $PDel^\omega$ is HO.

Proof. Obviously, we want to apply Lemma 2. Then we first need to show that our PDeIs contain c_{tot} .

- If $PDel_1$ and $PDel_2$ contain c_{tot} , then $PDel_1 \cup PDel_2$ trivially contains it too.
- If $PDel_1$ and $PDel_2$ contain c_{tot} , then $PDel_1 \otimes PDel_2$ contains $c_{tot} \otimes c_{tot} = c_{tot}$.
- If $PDel_1$ and $PDel_2$ contain c_{tot} , then $PDel_1 \rightsquigarrow PDel_2 \supseteq PDel_2$ contains it too.
- If $PDel$ contains c_{tot} , we can recreate c_{tot} by taking all $c_i = c_{tot}$ and whichever r_i . Thus, $PDel^\omega$ contains c_{tot} .

Next, we need to show that the $Nexts_f$ for the strategies corresponds to the generating sets in the theorem.

- We show $Nexts_{f_1 \cup f_2} = Nexts_{f_1} \cup Nexts_{f_2}$, and thus that $PHO_{f_1 \cup f_2}(PDel_1 \cup PDel_2) = HOProd(Nexts_{f_1 \cup f_2}) = HOProd(Nexts_{f_1} \cup Nexts_{f_2})$
 - Let $n \in Nexts_{f_1 \cup f_2}$. Then $\exists q \in f_1 \cup f_2 : obliv(q) = n$. By definition of union, $q \in f_1$ or $q \in f_2$. We fix $q \in f_1$ (the case $q \in f_2$ is symmetric). Then $n \in Nexts_{f_1}$.
We conclude that $n \in Nexts_{f_1} \cup Nexts_{f_2}$.
 - Let $n \in Nexts_{f_1} \cup Nexts_{f_2}$. We fix $n \in Nexts_{f_1}$ (as always, the other case is symmetric). Then $\exists q \in f_1 : obliv(q) = n$. As $q \in f_1$ implies $q \in f_1 \cup f_2$, we conclude that $n \in Nexts_{f_1 \cup f_2}$.
- We show $Nexts_{f_1 \otimes f_2} = \{n_1 \cap n_2 \mid n_1 \in Nexts_{f_1} \wedge n_2 \in Nexts_{f_2}\}$.
 - Let $n \in Nexts_{f_1 \otimes f_2}$. Then $\exists q \in f_1 \otimes f_2 : obliv(q) = n$. By definition of combination, $\exists q_1 \in f_1, \exists q_2 \in f_2 : q_1.round = q_2.round = q.round \wedge q = q_1 \otimes q_2$. This means $n = obliv(q) = obliv(q_1) \cap obliv(q_2)$.
We conclude that $n \in \{n_1 \cap n_2 \mid n_1 \in Nexts_{f_1} \wedge n_2 \in Nexts_{f_2}\}$.
 - Let $n \in \{n_1 \cap n_2 \mid n_1 \in Nexts_{f_1} \wedge n_2 \in Nexts_{f_2}\}$. Then $\exists n_1 \in Nexts_{f_1}, \exists n_2 \in Nexts_{f_2} : n = n_1 \cap n_2$. Because f_1 and f_2 are oblivious strategies, we can find $q_1 \in f_1$ such that $obliv(q_1) = n_1$, $q_2 \in f_2$ such that $obliv(q_2) = n_2$, and $q_1.round = q_2.round$.
Then $q = q_1 \otimes q_2$ is a state of $f_1 \otimes f_2$. We have $obliv(q) = n_1 \cap n_2 = n$.
We conclude that $n \in Nexts_{f_1 \otimes f_2}$.
- Trivially, $Nexts_f = Nexts_f$.

B.4 Domination by an Oblivious Strategy

To prove Theorem 5, we first show that the condition implies the domination by an oblivious strategy.

Lemma 7 (Sufficient Condition to be Dominated by an Oblivious Strategy). *Let $PDel$ be a delivered predicate. If*

- **(Total collection)** $PDel$ contains the total collection c_{tot} ,
- **(Symmetry up to a round)** $\forall c \in PDel, \forall p \in \Pi, \forall r > 0, \forall r' > 0, \exists c' \in PDel : c'[1, r' - 1] = c_{tot}[1, r' - 1] \wedge \forall q \in \Pi : c'(r', q) = c(r, p)$

then $PDel$ is dominated by an oblivious strategy.

Proof. Proved in [12, Thm 24].

Theorem ((5) Domination by Oblivious for Operations). *Let $PDel, PDel_1, PDel_2$ be delivered predicates that satisfy:*

- **(Total collection)** They contains the total collection c_{tot} ,
- **(Symmetry up to a round)** $\forall c$ a collection in the predicate, $\forall p \in \Pi, \forall r > 0, \forall r' > 0, \exists c'$ a collection in the predicate: $c'[1, r' - 1] = c_{tot}[1, r' - 1] \wedge \forall q \in \Pi : c'(r', q) = c(r, p)$

Then $PDel_1 \cup PDel_2, PDel_1 \otimes PDel_2, PDel_1 \rightsquigarrow PDel_2, PDel^\omega$ satisfy the same two conditions and are dominated by oblivious strategies.

Proof (Proof idea). Thanks to Lemma 7, we only have to show that the condition is maintained by the operations; the domination by an oblivious strategy follows directly.

For containing c_{tot} : $c_{tot} \cup c_{tot} = c_{tot}$; $c_{tot} \otimes c_{tot} = c_{tot}$; $c_{tot} \rightsquigarrow c_{tot} = c_{tot}$; and the succession of c_{tot} with itself again and again gives c_{tot} .

As for symmetry up to a round, we show its invariance. Let $p \in \Pi, r > 0$ and $r' > 0$.

- If $c \in PDel_1 \cup PDel_2$, then $c \in PDel_1 \vee c \in PDel_2$. We can then apply the condition for one of them to get c' .
- If $c \in PDel_1 \otimes PDel_2$, then $\exists c_1 \in PDel_1, \exists c_2 \in PDel_2 : c = c_1 \otimes c_2$. Applying the condition for c_1 and c_2 gives us c'_1 and c'_2 , and $c' = c'_1 \otimes c'_2$ satisfies the condition for c .
- If $c \in PDel_1 \rightsquigarrow PDel_2$, then $\exists c_1 \in PDel_1, \exists c_2 \in PDel_2, \exists r_{change} \geq 0 : c = c_1[1, r_{change}].c_2$. Applying the condition for c_1 at r and r' and for c_2 at $r - r_{change}$ and $r' - r_{change}$ gives us c'_1 and c'_2 , and $c' = c'_1[1, r_{change}].c'_2$ satisfies the condition for c .
- If $c \in PDel^\omega$, then $\exists (c_i)_{i \in \mathbb{N}^*}, \exists (r_i)_{i \in \mathbb{N}^*}$, the collections and indices defining c . Then let i the integer such that $r \in [r_i + 1, r_{i+1}]$. Applying the condition for $c_{i'}$ at $r - r_{i'}$ and $r' - r_{i'}$ with $i' \leq i$ gives us $c'_{i'}$, and $c' = c'_1[1, r_2 - r_1] \cdots c'_i[1, r_{i+1} - r_i].c_{i+1}[1, r_{i+2} - r_{i+1}] \cdots$ satisfies the condition for c .

Proof. Thanks to Lemma 7, we only have to show that the condition is maintained by the operations; the domination by an oblivious strategy follows directly.

We first prove that c_{tot} is still in the results of the operations.

- If $PDel_1$ and $PDel_2$ contain c_{tot} , then $PDel_1 \cup PDel_2$ trivially contains it too.
- If $PDel_1$ and $PDel_2$ contain c_{tot} , then $PDel_1 \otimes PDel_2$ contains $c_{tot} \otimes c_{tot} = c_{tot}$.
- If $PDel_1$ and $PDel_2$ contain c_{tot} , then $PDel_1 \rightsquigarrow PDel_2 \supseteq PDel_2$ contains it too.
- If $PDel$ contains c_{tot} , we can recreate c_{tot} by taking all $c_i = c_{tot}$ and whichever r_i . Thus, $PDel^\omega$ contains c_{tot} .

Then we show the invariance of the symmetry up to a round.

- Let $c \in PDel_1 \cup PDel_2$. Thus $c \in PDel_1$ or $c \in PDel_2$. We fix $c \in PDel_1$ (the other case is symmetric). Then for $p \in \Pi, r > 0$ and $r' > 0$, we get a $c' \in PDel_1$. satisfying the condition. And since $PDel_1 \subseteq PDel_1 \cup PDel_2$, we get $c' \in PDel_1 \cup PDel_2$.

We conclude that the condition still holds for $PDel_1 \cup PDel_2$.

- Let $c \in PDel_1 \otimes PDel_2$. Then $\exists c_1 \in PDel_1, \exists c_2 \in PDel_2 : c = c_1 \otimes c_2$. For $p \in \Pi, r > 0$ and $r' > 0$, our hypothesis on $PDel_1$ and $PDel_2$ ensures that there are $c'_1 \in PDel_1$ satisfying the condition for c_1 and $c'_2 \in PDel_2$ satisfying the condition for c_2 .

We argue that $c' = c'_1 \otimes c'_2$ satisfies the condition for c . Indeed, $\forall r'' < r', \forall q \in \Pi : c(r'', q) = c'_1(r'', q) \otimes c'_2(r'', q) = \Pi$ and $\forall q \in \Pi : c(r', q) = c'_1(r', q) \otimes c'_2(r', q) = c_1(r, p) \otimes c_2(r, p) = c(r, p)$.

We conclude that the condition still holds for $PDel_1 \otimes PDel_2$.

- Let $c \in PDel_1 \rightsquigarrow PDel_2$. Since if $c \in PDel_2$, the condition trivially holds by hypothesis, we study the case where succession actually happens. Hence, $\exists c_1 \in PDel_1, \exists c_2 \in PDel_2, \exists r_{change} > 0 : c = c_1[1, r_{change}].c_2$. For $p \in \Pi, r > 0$ and $r' > 0$, we separate two cases.

- if $r \leq r_{change}$, then our hypothesis on $PDel_1$ ensures that there is $c'_1 \in PDel_1$ satisfying the condition for c_1 . We argue that $c' = c'_1[1, r'].c_2 \in PDel_1 \rightsquigarrow PDel_2$ satisfies the condition for c .
Indeed, $\forall r'' < r', \forall q \in \Pi : c'(r'', q) = c'_1(r'', q) = \Pi$, and $\forall q \in \Pi : c'(r', q) = c_1(r, p) = c(r, p)$
- if $r > r_{change}$, then our hypothesis on $PDel_2$ ensures that there is $c'_2 \in PDel_2$ satisfying the condition for c_2 at p and $r - r_{change}$. That is, $c'_2[1, r' - 1] = c_{tot}[1, r' - 1] \wedge \forall q \in \Pi : c'_2(r', q) = c_2(r - r_{change}, p)$ We argue that $c' = c'_2 \in PDel_1 \rightsquigarrow PDel_2$ satisfies the condition for c .
Indeed, $\forall r'' < r', \forall q \in \Pi : c'_2(r'', q) = \Pi$, and $\forall q \in \Pi : c'_2(r', q) = c_2(r - r_{change}, p) = c(r, p)$

We conclude that the condition still holds for $PDel_1 \rightsquigarrow PDel_2$.

- Let $c \in PDel^\omega$. Let (c_i) and (r_i) be the collections and indices defining c . We take $p \in \Pi, r > 0$ and $r' > 0$. Let $i > 0$ be the integer such that $r \in [r_i + 1, r_{i+1}]$. By hypothesis on $PDel$, There is $c'_i \in PDel$ satisfying the condition for c_i at p and $r - r_i$. That is, $c'_i[1, r' - 1] = c_{tot}[1, r' - 1] \wedge \forall q \in \Pi : c'_i(r', q) = c_i(r - r_i, p)$.
We argue that $c'_i \in PDel$ satisfies the condition for c . Indeed, $\forall r'' \leq r', \forall q \in \Pi$, we have: $c'_i(r'', q) = \Pi$ and $\forall q \in \Pi : c'_i(r', q) = c_i(r - r_i, p) = c(r, p)$.
We conclude that the condition still holds for $PDel^\omega$.

C Proofs for Conservative Strategies

C.1 Minimal Conservative Strategies

We use a necessary and sufficient condition for an oblivious strategy to be valid in the rest of the proofs.

Lemma 8 (Necessary and Sufficient Condition for Validity of a Conservative Strategy). *Let $PDel$ be a delivered predicate and f be a conservative strategy. Then f is valid for $PDel \iff f \supseteq \{q \in Q \mid \exists c \in PDel, \exists p \in \Pi, \forall r \leq q.\text{round} : q(r) = c(r, p)\}$.*

Proof. From the version in [12], f has to satisfy $\forall CDel \in PDel, \forall r > 0, \forall j \in \Pi : \langle r, \{\langle r', k \mid r' \leq r \wedge k \in CDel(r', j) \} \rangle \in Nexts_f^R$.

We show the equivalence of this condition with our own, which allow us to conclude by transitivity of equivalence.

- (\implies) assume our condition holds and prove the one from [12],
Let $c \in PDel, r > 0$ and $p \in \Pi$: we want to show that $q = \langle r, \{\langle r', k \mid r' \leq r \wedge k \in c(r', p) \} \rangle \in Nexts_f^R$. That is to say, that all states whose past and present correspond to this conservative state are accepted by f . Let q' such that $cons(q') = q$, that is $q'.\text{round} = q.\text{round} = r$ and $\forall r' \leq r : q'(r') = q(r') = c(r', p)$. We have the collection and the round to apply our condition, and thus $q \in f$.
- (\impliedby) assume the condition from [12] holds and we prove ours.
Let q such that $\exists c \in PDel, \exists p \in \Pi, \forall r \leq q.\text{round} : q(r) = c(r, p)$.
Then $cons(q) = \langle q.\text{round}, \{\langle r, k \mid r \leq q.\text{round} \wedge k \in c(r, p) \} \rangle$. This conservative state is in $Nexts_f^R$ by hypothesis.
We conclude that $q \in f$.

Lemma ((3) Domination of Minimal Conservative Strategy). *Let $PDel$ be a delivered predicate and f_{min} be its minimal conservative strategy. Then f_{min} dominates the conservative strategies for $PDel$.*

Proof. First, f_{min} is valid for $PDel$ by application of Lemma 8. Next, we take another conservative strategy f , valid for $PDel$. Lemma 8 gives us that $f_{min} \subseteq f$. Hence, when f_{min} allow a change of round, so does f . This entails that all executions of f_{min} for $PDel$ are also executions of f for $PDel$, and thus that the $PHO_{f_{min}}(PDel) \subseteq PHO_f(PDel)$.

C.2 Operations Maintain Minimal Conservative Strategy

Theorem ((6) Minimal Conservative Strategy for Union). *Let $PDel_1, PDel_2$ be two PDels, f_1 the minimal conservative strategy for $PDel_1$, and f_2 the minimal conservative strategy for $PDel_2$. Then $f_1 \cup f_2$ is the minimal conservative strategy for $PDel_1 \cup PDel_2$.*

Proof. We only have to show that $f_1 \cup f_2$ is equal to Definition 18.

- (\supseteq) Let q be a state such that $\exists c \in PDel_1 \cup PDel_2, \exists p \in \Pi$ such that $\forall r \leq q.\text{round} : q(r) = c(r, p)$. If $c \in PDel_1$, then $q \in f_1$, because f_1 is the minimal conservative strategy for $PDel_1$, and by application of Lemma 8. Thus, $q \in f_1 \cup f_2$. If $c \in PDel_2$, the same reasoning apply with f_2 in place of f_1 . We conclude that $q \in f_1 \cup f_2$.
- (\subseteq) Let $q \in f_1 \cup f_2$. This means that $q \in f_1 \vee q \in f_2$. The case where it is in both can be reduced to any of the two. If $q \in f_1$, then by minimality of f_1 $\exists c_1 \in PDel_1, \exists p_1 \in \Pi$ such that $\forall r \leq q.\text{round} : q(r) = c_1(r, p_1)$. $PDel_1 \subseteq PDel_1 \cup PDel_2$, thus $c_1 \in PDel_1 \cup PDel_2$. We found the c and p necessary to show q is in the minimal conservative strategy for $PDel_1 \cup PDel_2$. If $q \in f_2$, the reasoning is similar to the previous case, replacing f_1 by f_2 and $PDel_1$ by $PDel_2$.

Theorem ((7) Minimal Conservative Strategy for Combination). *Let $PDel_1, PDel_2$ be two symmetric PDels, f_1 the minimal conservative strategy for $PDel_1$, and f_2 the minimal conservative strategy for $PDel_2$. Then $f_1 \otimes f_2$ is the minimal conservative strategy for $PDel_1 \otimes PDel_2$.*

Proof. We only have to show that $f_1 \otimes f_2$ is equal to Definition 18.

- (\supseteq) Let q be a state such that $\exists c \in PDel_1 \otimes PDel_2, \exists p \in \Pi$ such that $\forall r \leq q.\text{round} : q(r) = c(r, p)$. By definition of c , $\exists c_1 \in PDel_1, \exists c_2 \in PDel_2 : c_1 \otimes c_2 = c$. We take q_1 such that $q_1.\text{round} = q.\text{round}$ and $\forall r > 0 :$

$$\begin{pmatrix} q_1(r) = c_1(r, p) & \text{if } r \leq q.\text{round} \\ q_1(r) = q(r) & \text{otherwise} \end{pmatrix}$$
 We also take q_2 such that $q_2.\text{round} = q.\text{round}$ and $\forall r > 0 :$

$$\begin{pmatrix} q_2(r) = c_2(r, p) & \text{if } r \leq q.\text{round} \\ q_2(r) = q(r) & \text{otherwise} \end{pmatrix}$$

Then by validity of f_1 and f_2 (since they are minimal conservative strategies) and by application of Lemma 8, we get $q_1 \in f_1$ and $q_2 \in f_2$. We also see that $q = q_1 \otimes q_2$. Indeed, for $r \leq q.\text{round}$, we have $q(r) = c(r, p) = c_1(r, p) \cap c_2(r, p) = q_1(r) \cap q_2(r)$; and for $r > q.\text{round}$, we have $q(r) = q(r) \cap q(r) = q_1(r) \cap q_2(r)$.

Therefore $q \in PDel_1 \otimes PDel_2$.

- (\subseteq) Let $q \in f_1 \otimes f_2$. By definition of $f_1 \otimes f_2$, $\exists q_1 \in f_1, \exists q_2 \in f_2$ such that $q_1.\text{round} = q_2.\text{round} = q.\text{round}$ and $q = q_1 \otimes q_2$. Since f_1 and f_2 are minimal conservative strategies of their respective PDels, $\exists c_1 \in PDel_1, \exists p_1 \in \Pi$ such that $\forall r \leq q.\text{round} : q_1(r) = c_1(r, p_1)$; and $\exists c_2 \in PDel_2, \exists p_2 \in \Pi$ such that $\forall r \leq q.\text{round} : q_2(r) = c_2(r, p_2)$.

By symmetry of $PDel_2$, $\exists c'_2 \in PDel_2$ such that $\forall r \leq q.\text{round} : c'_2(r, p_1) = c_2(r, p_2)$. Hence, $\forall r \leq q.\text{round} : q_2(r) = c'_2(r, p_1)$.

By taking $c = c_1 \otimes c_2$, we get $\forall r \leq q.\text{round} : q(r) = q_1(r) \cap q_2(r) = c_1(r, p_1) \cap c_2(r, p_1) = c(r, p_1)$.

We found c and p showing that q is in the minimal conservative strategy for $PDel_1 \otimes PDel_2$.

Theorem ((8) Minimal Conservative Strategy for Succession). *Let $PDel_1, PDel_2$ be two symmetric PDels, f_1 the minimal conservative strategy for $PDel_1$, and f_2 the minimal conservative strategy for $PDel_2$. Then $f_1 \rightsquigarrow f_2$ is the minimal conservative strategy for $PDel_1 \rightsquigarrow PDel_2$.*

Proof. We only have to show that $f_1 \rightsquigarrow f_2$ is equal to Definition 18.

- (\supseteq) Let q be a state such that $\exists c \in PDel_1 \rightsquigarrow PDel_2, \exists p \in \Pi$ such that $\forall r' \leq q.\text{round} : q(r') = c(r', p)$. By definition of c , $\exists c_1 \in PDel_1, \exists c_2 \in PDel_2, \exists r > 0 : c = c_1[1, r].c_2$.
 - If $r = 0$, then $c[1, r] = c_2[1, r]$, and thus $\forall r' \leq q.\text{round} : q(r') = c_2(r', p)$. The validity of f_2 and Lemma 8 then allow us to conclude that $q \in f_2$ and thus that $q \in f_1 \rightsquigarrow f_2$.
 - If $r > 0$, we have two cases to consider.
 - * If $q.\text{round} \leq r$, then $\forall r' \leq q.\text{round} : q(r') = c_1(r', p)$. We conclude by f_1 and application of Lemma 8 that $q \in f_1$ and thus that $q \in f_1 \rightsquigarrow f_2$.
 - * If $q.\text{round} > r$, then $c[1, q.\text{round}] = c_1[1, r].c_2[1, q.\text{round} - r]$. We take q_1 such that $q_1.\text{round} = r$ and $\forall r' > 0 :$

$$\begin{pmatrix} q_1(r') = c_1(r', p) & \text{if } r' \leq q_1.\text{round} \\ q_1(r') = q(r') & \text{otherwise} \end{pmatrix}$$
. We also take q_2 such that $q_2.\text{round} = q.\text{round} - r$ and $\forall r' > 0 :$

$$\begin{pmatrix} q_2(r') = c_2(r', p) & \text{if } r' \leq q_2.\text{round} \\ q_2(r') = q(r' - q.\text{round}) & \text{otherwise} \end{pmatrix}$$
. Then by validity of f_1 and f_2 , and by application of Lemme 8, we get $q_1 \in f_1$ and $q_2 \in f_2$. We also see that $q = q_1 \rightsquigarrow q_2$. Indeed, for $r' \leq q_1.\text{round} = r$, we have $q(r') = c(r', p) = c_1(r', p) = q_1(r')$; for $r' \in [q_1.\text{round} + 1, q.\text{round}]$, we have $q(r') = c(r', p) = c_2(r' - r, p) = q_2(r' - r)$ and for $r' > q.\text{round}$ we have $q(r') = q_2(r' - q.\text{round})$. We conclude that $q \in f_1 \rightsquigarrow f_2$.
- (\subseteq) Let $q \in f_1 \rightsquigarrow f_2$. By definition of succession for strategies, there are three possibilities for q .
 - If $q \in f_1$, then by minimality of f_1 $\exists c_1 \in PDel_1, \exists p_1 \in \Pi : \forall r \leq q.\text{round} : q(r) = c_1(r, p_1)$. Let $c_2 \in PDel_2$. We take $c = c_1[1, q.\text{round}].c_2$; we have $c \in c_1 \rightsquigarrow c_2$. Then, $\forall r \leq q.\text{round} : q(r) = c_1(r, p_1) = c(r, p_1)$. We found c and p showing that q is in the minimal conservative strategy for $PDel_1 \rightsquigarrow PDel_2$.
 - If $q \in f_2$, then by minimality of f_2 $\exists c_2 \in PDel_2, \exists p_2 \in \Pi : \forall r \leq q.\text{round} : q(r) = c_2(r, p_2)$. As $PDel_2 \subseteq PDel_1 \rightsquigarrow PDel_2$, thus $c_2 \in PDel_1 \rightsquigarrow PDel_2$.

We found c and p showing that q is in the minimal conservative strategy for $PDel_1 \rightsquigarrow PDel_2$.

- There are $q_1 \in f_1$ and $q_2 \in f_2$ such that $q = q_1 \rightsquigarrow q_2$.

Because f_1 and f_2 are the minimal conservative strategies of their respective PDels, $\exists c_1 \in PDel_1, \exists p_1 \in \Pi$ such that $\forall r \leq q_1.\text{round} : q_1(r) = c_1(r, p_1)$; and $\exists c_2 \in PDel_2, \exists p_2 \in \Pi$ such that $\forall r \leq q_2.\text{round} : q_2(r) = c_2(r, p_2)$.

By symmetry of $PDel_2$, $\exists c'_2 \in PDel_2 : \forall r \leq q_2.\text{round} : c'_2(r, p_1) = c_2(r, p_2)$. Hence, $\forall r \leq q_2.\text{round} : q_2(r) = c'_2(r, p_1)$.

By taking $c = c_1[1, q_1.\text{round}].c'_2$, we have $c \in c_1 \rightsquigarrow c'_2$. Then $\forall r \leq q.\text{round} = q_1.\text{round} + q_2.\text{round} :$

$$\left(\begin{array}{l} q(r) = q_1(r) \\ \quad = c_1(r, p_1) \quad \text{if } r \leq q_1.\text{round} \\ \quad = c(r, p_1) \\ q(r) = q_2(r - q_1.\text{round}) \\ \quad = c'_2(r - q_1.\text{round}, p_1) \text{ if } r \in [q_1.\text{round} + 1, q_1.\text{round} + q_2.\text{round}] \\ \quad = c(r, p_1) \end{array} \right).$$

We found c and p showing that q is in the minimal conservative strategy for $PDel_1 \rightsquigarrow PDel_2$.

Theorem ((9) Minimal Conservative Strategy for Repetition). *Let $PDel$ be a symmetric PDel, and f be its minimal conservative strategy. Then f^ω is the minimal conservative strategy for $PDel^\omega$.*

Proof. We only have to show that f^ω is equal to Definition 18.

- (\supseteq) Let q be a state such that $\exists c \in PDel^\omega, \exists p \in \Pi$ such that $\forall r \leq q.\text{round} : q(r) = c(r, p)$. By definition of repetition, $\exists (c_i)_{i \in \mathbb{N}^*}, \exists (r_i)_{i \in \mathbb{N}^*}$ such that $r_1 = 0$ and $\forall i \in \mathbb{N}^* : (c_i \in PDel \wedge r_i < r_{i+1} \wedge c[r_i + 1, r_{i+1}] = c_i[1, r_{i+1} - r_i])$. Let k be the biggest integer such that $r_k \leq q.\text{round}$. We consider two cases.

- If $r_k = q.\text{round}$, then $c[1, r] = c_1[1, r_2 - r_1].c_2[1, r_3 - r_2] \dots c_{k-1}[1, r_k - r_{k-1}]$. We take for $i \in [1, k-1] : q_i$ the state such that $q_i.\text{round} = r_{i+1} - r_i$ and $\forall r > 0 :$

$$\left(\begin{array}{l} q_i(r) = c_i(r, p) \quad \text{if } r \leq q_i.\text{round} \\ q_i(r) = q(r + \sum_{j \in [1, i-1]} q_j.\text{round}) \text{ otherwise} \end{array} \right).$$

By validity of f and by application of Lemma 8, for $i \in [1, k-1]$ we have $q_i \in f$. We see that $\forall r > 0 : q(r) = (q_1 \rightsquigarrow \dots \rightsquigarrow q_{k-1})(r)$. Indeed, $\forall r \in [r_i + 1, r_{i+1}] : q(r) = c(r, p) = c_i(r - r_i, p) = q_i(r - r_i)$; and for $r > q.\text{round} : q(r) = q_{k-1}(r - \sum_{j \in [1, k-1]} q_j.\text{round})$.

We conclude that $q \in f^\omega$.

- If $q.\text{round} > r_k$, we can apply the same reasoning as in the previous case, the only difference being $c[1, r] = c_1[1, r_2 - r_1].c_2[1, r_3 - r_2] \dots c_{k-1}[1, r_k - r_{k-1}].c_k[1, r - r_k]$.
- (\subseteq) Let $q \in f^\omega$. By definition of f^ω , $\exists q_1, q_2, \dots, q_k \in f : q = q_1 \rightsquigarrow q_2 \rightsquigarrow \dots \rightsquigarrow q_k$.

Since f is the minimal conservative strategy of $PDel$, $\exists c_1, c_2, \dots, c_k \in PDel, \exists p_1, p_2, \dots, p_k \in \Pi : \forall i \in [1, k] q_i = \langle q_i.\text{round}, \{ \langle r, j \rangle \mid r \leq q_i.\text{round} \wedge j \in c_i(r, p_i) \} \rangle$.

By symmetry of $PDel$, for all $i \in [2, k]$, $\exists c'_i \in PDel, \forall r \leq q_i.\text{round} : c'_i(r, p_1) = c_i(r, p_i)$.

We take $c = c_1[1, q_1.\text{round}].c'_2[1, q_2.\text{round}]\dots c'_{k-1}[1, q_{k-1}.\text{round}].c'_k$, thus $c \in c_1 \rightsquigarrow c'_2 \rightsquigarrow \dots \rightsquigarrow c'_k$. Then $\forall r \leq q.\text{round} = \sum_{i \in [1, k]} q_i.\text{round}$, if $r \in [\sum_{i \in [1, i-1]} q_i.\text{round} +$

$$1, \sum_{i \in [1, i]} q_i.\text{round}], \text{ we have } \begin{pmatrix} q(r) = q_i(r - \sum_{i \in [1, i-1]} q_i.\text{round}) \\ = c_i(r - \sum_{i \in [1, i-1]} q_i.\text{round}, p_1) \\ = c(r, p_1) \end{pmatrix}.$$

We found c and p showing that q is in the minimal conservative strategy for $PDel^\omega$.

C.3 Computing Heard-Of Predicates

Theorem ((10) Upper Bounds on HO of Minimal Conservative Strategies). *Let $PDel, PDel_1, PDel_2$ be PDels containing c_{tot} . Let $f^{cons}, f_1^{cons}, f_2^{cons}$ be their respective minimal conservative strategies, and $f^{obliv}, f_1^{obliv}, f_2^{obliv}$ be their respective minimal oblivious strategies. Then:*

- $PHO_{f_1^{cons} \cup f_2^{cons}}(PDel_1 \cup PDel_2) \subseteq HOProd(Nexts_{f_1^{obliv}} \cup Nexts_{f_2^{obliv}})$.
- $PHO_{f_1^{cons} \rightsquigarrow f_2^{cons}}(PDel_1 \rightsquigarrow PDel_2) \subseteq HOProd(Nexts_{f_1^{obliv}} \cup Nexts_{f_2^{obliv}})$.
- $PHO_{f_1^{cons} \otimes f_2^{cons}}(PDel_1 \otimes PDel_2) \subseteq HOProd(\{n_1 \cap n_2 \mid n_1 \in Nexts_{f_1^{obliv}} \wedge n_2 \in Nexts_{f_2^{obliv}}\})$.
- $PHO_{(f^{cons})^\omega}(PDel^\omega) \subseteq HOProd(Nexts_{f^{obliv}})$.

Proof. A oblivious strategy is a conservative strategy. Therefore, the minimal conservative strategy always dominates the minimal oblivious strategy. Hence, we get an upper bound on the heard-of predicate of the minimal conservative strategies by applying Theorem 4.