

Assuming Just Enough Fairness to make Session Types Complete for Lock-freedom

Rob van Glabbeek

Data61, CSIRO and UNSW, Australia

Sydney, Australia

Email: rvg@cs.stanford.edu

Peter Höfner

Australian National University

Canberra, Australia

Email: peter.hoefner@anu.edu.au

Ross Horne

Computer Science, University of Luxembourg

Esch-sur-Alzette, Luxembourg

Email: ross.horne@uni.lu

Abstract—We investigate how different fairness assumptions affect results concerning *lock-freedom*, a typical liveness property targeted by session type systems. We fix a minimal session calculus and systematically take into account all known fairness assumptions, thereby identifying precisely three interesting and semantically distinct notions of lock-freedom, all of which having a sound session type system. We then show that, by using a general merge operator in an otherwise standard approach to global session types, we obtain a session type system complete for the strongest amongst those notions of lock-freedom, which assumes only *justness* of execution paths, a minimal fairness assumption for concurrent systems.

I. INTRODUCTION

It has long been known that there is an intimate relationship between liveness properties and fairness assumptions. Seminal work by Owicki and Lamport [1] draws attention to the fact that liveness properties, such as “each request will eventually be answered” are indispensable to create correct concurrent programs.

Typically, a liveness property does not hold for all execution paths of a concurrent system: imagine two sellers and two buyers: *buyer1* repeatedly requests product *A* from *seller1*, who is able to sell the product. Similarly, *buyer2* and *seller2* are able to exchange product *B*. Assuming that both buyers try to request infinitely many products, there is an infinite execution path where product *A* is always requested and bought, and *B* is never sold. When taking all infinite execution paths into consideration, the rudimentary liveness property mentioned by Owicki and Lamport does not hold. Ranging over all infinite or completed executions – the default assumption for many model checkers – essentially assumes only that the system as a whole progresses if there is some work to do and there is no deadlock.

When reasoning about starvation-sensitive liveness properties, i.e. properties that avoid situations where a component wants to do something but is denied forever, Owicki and Lamport state explicitly that such liveness properties depend on a fairness assumption.

Assuming that the parties in our example act independently, claiming that the aforementioned liveness property fails is unrealistic, for both sellers should be able to react on any request. It is reasonable to make some fairness assumption

that ensures that the parties requesting and selling *A* do not impair the parties involved with *B*. This simple example can be used as a litmus test that any realistic fairness assumption for a concurrent system should pass.

Thus, liveness properties have to be parametrised with a fairness assumption that rules out potential executions of a system. As the fairness assumption becomes weaker (permitting more executions), the liveness property becomes stronger (systems can do more, so the liveness property is more likely to be rejected).

A reason why there exist different notions of fairness is that some notions are not realistic for some applications. For example, an implication of making the strongest of all fairness assumptions might be that you will phone everyone in your phone book repeatedly, which is unlikely. The minimal assumption *justness* [2] does not entail this, but it does imply that you will not be prevented from having a phone conversation due to unrelated calls between others. A recent survey [2] of fairness assumptions classifies dozens of semantically distinct notions by their strength in ruling out potential executions. Thus, for every liveness property, there are dozens of incarnations of that property obtained simply by varying the underlying fairness assumption.

Not all liveness properties obtained by varying fairness assumptions are semantically distinct. We identify two key reasons why liveness properties coincide: the (fixed) choice of process model and the choice of liveness property.

In this paper, we fix the process model to be a core synchronous session calculus featuring an internal and external choice [3], [4], which is frequently studied in the context of session types. We also fix the liveness properties to follow a scheme for *lock-freedom* [5], [6], which has emerged as one of the most important liveness properties for multiparty session calculi and related calculi, such as the linear π -calculus. Lock-freedom is essentially the absence of starvation, as described above. Clearly, the choice of the fairness assumption will influence whether a system is lock-free.

The restriction to session calculi, for which *session type systems* exist, allows us to answer the following question:

For a given fairness assumption, does there exist a session type system that is sound and/or complete, in the sense that a network is lock-free if and/or only if it is well-typed?

$$\begin{array}{c}
(\mathbb{N}_1 \parallel \mathbb{N}_2) \parallel \mathbb{N}_3 \equiv \mathbb{N}_1 \parallel (\mathbb{N}_2 \parallel \mathbb{N}_3) \quad \mathbb{M} \parallel \mathbb{N} \equiv \mathbb{N} \parallel \mathbb{M} \quad \mathbb{M} \parallel 0 \equiv \mathbb{M} \\
\\
\frac{\mathbb{N} \equiv \mathbb{N}' \quad \mathbb{N}' \xrightarrow{\alpha} \mathbb{M}' \quad \mathbb{M}' \equiv \mathbb{M}}{\mathbb{N} \xrightarrow{\alpha} \mathbb{M}} \quad \frac{k \in I}{p[\bigoplus_{i \in I} p_i! \lambda_i; \mathbb{T}_i] \parallel \mathbb{N} \xrightarrow{\tau} p[\ulcorner p_i! \lambda_i; \mathbb{T}_i \urcorner] \parallel \mathbb{N}} \\
\\
\frac{}{p[\mu X.T] \parallel \mathbb{N} \xrightarrow{\tau} p[\mathbb{T}\{\mu X.T/X\}] \parallel \mathbb{N}} \quad \frac{k \in I}{p_k[\ulcorner q! \lambda_k; \mathbb{U} \urcorner] \parallel q[\sum_{i \in I} p_i? \lambda_i; \mathbb{T}_i] \parallel \mathbb{N} \xrightarrow{p_k \rightarrow q: \lambda_k} p_k[\mathbb{U}] \parallel q[\mathbb{T}_k] \parallel \mathbb{N}}
\end{array}$$

Fig. 1. The default semantics for networks that we fix for this study.

Our systematic study yields the following main contributions.

- 1) We classify the notions of lock-freedom that arise by taking every notion of fairness in the survey [2] and using them to instantiate a parameter in a general scheme for lock-freedom. The resulting classification includes classic notions of lock-freedom of session calculi found in the literature. Hence it relates these notions as well. However, we discover that the notion of lock-freedom which arises from *justness* is new to the literature.
- 2) We introduce a generalisation of the projection mechanism of global types onto threads, which uses the most general possible merge operator. This solves the problem that session type systems employing global types without an explicit parallel composition operator are incomplete, in the sense that there are lock-free networks that cannot be typed. This leads to the following main result.
- 3) We prove that our session type system is complete for lock-freedom, when assuming justness. To the best of our knowledge, this is the first completeness result of this kind. We delineate the scope of our completeness result by showing that completeness does not hold for weaker notions of lock-freedom.
- 4) We prove that more notions of lock-freedom coincide when restricting to race-free networks. Furthermore, race-free networks are sound for all notions of lock-freedom, whenever we assume at least *justness*.

Following [7], [8], we employ session types that abstract from the concrete types (e.g. `Bool` or `Nat`) of messages, using labels λ instead. As a result, systems and types have a fairly similar syntax. It is fairly trivial to move from our session type system with labels to one with data and data types.

Structure of the paper: Section II introduces our session calculus and a spectrum of fairness assumptions, and then systematically classifies the resulting spectrum of lock-freedom properties. Section III presents our session type system featuring a general merge operator and guarded types, which we prove to be complete with respect to $\mathcal{L}(J)$ – the notion of lock-freedom arising from the assumption of justness – for all networks. Section IV considers race-free networks in order to explore the scope of soundness results. Section V situates our results with respect to notions of lock-freedom from the literature.

II. THE SCOPE: A SESSION CALCULUS, ITS KEY FAIRNESS NOTIONS AND LIVENESS PROPERTIES

In this section, we define the session calculus and a scheme for lock freedom. We also explain various fairness assumptions and illustrate their differences through separating examples.

A. Syntax and semantics for threads and networks

Our session calculus features finitely many recursive *threads* that send and receive messages. Threads, uniquely identified by location names, feature an internal choice $\bigoplus p_i! \lambda_i$ between messages labelled λ_i sent to locations p_i (a choice made at run-time entirely by the sending thread), and an external choice $\sum p_i? \lambda_i$ amongst messages received (meaning that the thread is ready to receive different messages λ_i from p_i , but cannot influence which of them will eventually come through).

$$\begin{array}{l|l}
\mathbb{T} := & \text{OK} \\
& \bigoplus_{i \in I} p_i! \lambda_i; \mathbb{T}_i \\
& \sum_{i \in I} p_i? \lambda_i; \mathbb{T}_i \\
& X \\
& \mu X.T \\
\mathbb{N} := & p[\mathbb{T}] \\
& 0 \\
& \mathbb{N} \parallel \mathbb{N}
\end{array}$$

The index sets I are finite, and in the case of $\bigoplus_{i \in I}$ also non-empty. We enforce guarded recursion by excluding threads of the form $\mu X.X$ or $\mu X.\mu Y.T$. If $p[\mathbb{T}]$ is a sub-expression of a network \mathbb{N} , then p is called a *location* of \mathbb{N} . In a network \mathbb{N} , all locations are required to be distinct and all threads closed, meaning that each occurrence of a variable X is in the scope of a recursion $\mu X.T$. Moreover, in each sub-expression $p_k! \lambda_k$ or $p_k? \lambda_k$, the p_k must be a location of \mathbb{N} . We may elide `OK`; we write $p_1! \lambda_1; \mathbb{T}_1 \oplus \dots \oplus p_n! \lambda_n; \mathbb{T}_n$ for $\bigoplus_{i \in \{1, \dots, n\}} p_i! \lambda_i; \mathbb{T}_i$, and $p_1? \lambda_1; \mathbb{T}_1 + \dots + p_n? \lambda_n; \mathbb{T}_n$ for $\sum_{i \in \{1, \dots, n\}} p_i? \lambda_i; \mathbb{T}_i$. In particular, we write $p? \lambda; \mathbb{T}$ in case I is a singleton set. We follow a recent trend allowing inputs in an external choice to listen to different locations [7], [9], which allows us to broaden the scope of our investigation.

A reduction semantics for our session calculus: The rules for our session calculus, presented in Figure 1, are fairly standard. In this semantics, an output that a thread has committed to can interact synchronously with some input in an external choice. Also, recursion is unfolded by a τ -transition and the standard associativity and commutativity of parallel composition can be applied to enable any transition.

A design decision, we will demonstrate to be significant, is that there is a τ -transition for resolving all internal choices.

To ensure that singleton internal choices perform only one τ -transition (and not a diverging sequence of τ -transitions), the transition ends in a *network state* that is not a syntactically valid network. *Network states* are comprised of located *thread states*, which due to the annotation \ulcorner , are not necessarily threads themselves.

B. Fairness notions for session calculi

We now discuss three fairness assumptions for our session calculus. A fairness assumption restricts the set of complete execution paths, here simply referred to as *paths*.

Definition 1: A path consists of a network state \mathbb{N}_0 and a maximal list of transitions $\mathbb{N}_i \xrightarrow{\alpha_i} \mathbb{N}_{i+1}$, permitted by Figure 1.

Maximality ensures that either the list is infinite or the final network state has no outgoing transition, that is, we restrict ourselves to *complete* execution paths.

A *fairness notion* \mathcal{F} characterises a subset of all paths as the *fair* ones, modelling executions that we assume can actually occur; we refer to such paths as *\mathcal{F} -fair paths*. It is required to satisfy the condition of *feasibility* [10], saying that each finite prefix of a path is also a prefix of a fair path. One notion of fairness \mathcal{G} is *stronger* than another one \mathcal{F} – in symbols $\mathcal{F} \preceq \mathcal{G}$ – if it rules out more paths as unfair.

A network \mathbb{N} *successfully terminates* under a fairness notion \mathcal{F} iff all fair paths successfully terminate, i.e., all components of \mathbb{N} eventually take the form $p[\text{OK}]$.

A liveness property, or more generally a *linear-time property*, is formalised as a property φ of paths. It holds for network state \mathbb{N}_0 under a certain fairness assumption iff all fair paths starting in \mathbb{N}_0 satisfy φ .

B.1 Strong and weak fairness: In [2], the concepts of *strong and weak fairness* are parametrised by the notion of a *task*. What a task is may differ from one notion of fairness to another, but for each task it should be clear when it is *enabled* in a network state, and when a path *engages* in a task. A task T is said to be *relentlessly enabled* on a path π if each suffix of π contains a network state in which T is enabled; it is *perpetually enabled* if it is enabled in all network states of π . A path π is *strongly fair* if, for each suffix π' of π , each task that is relentlessly enabled on π' is engaged in by π' . It is *weakly fair* if, for each suffix π' of π , each task that is perpetually enabled on π' is engaged in by π' .

Given a notion of a task t , the concept of strong fairness St is always stronger than its weak counterpart Wt , i.e., $\text{Wt} \preceq \text{St}$.

In [2], several notions of fairness found in the literature are characterised through formalising what constitutes a task. *Fairness of transitions* is obtained by taking the tasks to be the transitions. Such a task is enabled in a network state \mathbb{N} if \mathbb{N} is the source state of that transition. A path π engages in a transition if that transition occurs in π .

Fact 1: Strong fairness of transitions (ST) characterises exactly those paths π with the property that whenever a transition is relentlessly enabled on π then the transition must be taken infinitely often on π ; it rules out all other paths.

In [2], it is shown that for finite-state systems strong fairness of transitions (ST) is the strongest feasible notion of fairness.

Example 1: Consider the following network where a *buyer* chooses to talk to or to buy a product from a *seller*, after which the order is shipped.

$$\begin{aligned} & \text{buyer} \llbracket \mu X. (\text{seller!talk}; X \oplus \text{seller!buy}) \rrbracket \\ \parallel & \text{seller} \llbracket \mu Y. (\text{buyer?talk}; Y \\ & \quad + \text{buyer?buy}; \text{shipper!order}) \rrbracket \\ \parallel & \text{shipper} \llbracket \text{seller?order} \rrbracket \end{aligned}$$

The network successfully terminates when assuming ST, for in the only infinite execution the τ -transition belonging to instruction *seller!buy* is relentlessly enabled but never taken.

A notion of task that figures prominently in the literature is that of a *component*. A component is one of the prime elements in a parallel composition – in a network expression it is completely determined by its location. Each transition involves either one or two components. A component is *enabled* in a network state iff a transition involving that component is enabled; a path *engages* in a component iff it contains a transition that involves that component.

We define a function *comp* which returns for a transition the set of components participating in the transition. Each transition labelled τ involves exactly one component (location) evident from the rule; each transition labelled $p \rightarrow q: \lambda$ involves exactly two components, p and q . This defines strong and weak fairness of components.

Fact 2: Strong fairness of components (SC) characterises the paths π such that, for any location p , if there are transitions involving p relentlessly enabled on π , then a transition that involves p must be taken infinitely often on π .

Fact 3: A path π satisfies *weak fairness of components (WC)* whenever, for every location p , if some transition involving p is, from some state onwards, perpetually enabled, then a transition that involves p occurs infinitely often in π .

Under the fairness assumption SC, Example 1 does not successfully terminate, for there is an infinite path where, alternately, the buyer performs a τ -transition to select the left branch of its choice and then the *buyer* and *seller* talk to each other. Along this path there is never a transition enabled that involves the *shipper*; hence that branch need never be taken. This illustrates that SC allows strictly more paths than ST, i.e., $\text{SC} \not\preceq \text{ST}$.

Example 2: To see that SC excludes some paths, consider the following network.

$$\begin{aligned} & \text{seller} \llbracket \mu X. (\text{buyer1?order1}; X + \text{buyer2?order2}) \rrbracket \\ \parallel & \text{buyer1} \llbracket \mu Y. \text{seller!order1}; Y \rrbracket \\ \parallel & \text{buyer2} \llbracket \text{seller!order2} \rrbracket \end{aligned}$$

The above network terminates under SC (albeit in a state where *buyer1* has not successfully terminated), for, in any infinite execution, a transition from *buyer2* is relentlessly enabled but never taken. It does not need to terminate under weak

fairness of components, for no transition is enabled perpetually due to the τ -transitions that unfold the recursion after each communication.

Guaranteeing termination in this example seems wrong as the fairness assumption constrains the ‘free will’ of the *seller* in the sense that they have to sell items to *buyer2*. Therefore we will introduce a weaker fairness assumption.

B.2 Justness: We consider a minimal notion of fairness that guarantees only that concurrent transitions cannot prevent each other from happening. Informally, two transitions are concurrent if no component is involved in both transitions.

Definition 2: Two transitions t and u are *concurrent*, notation $t \smile u$, if $\text{comp}(t) \cap \text{comp}(u) = \emptyset$.

Justness guarantees that once a transition is enabled that stems from a set of parallel components, one (or more) of these components will eventually partake in a transition.

Definition 3: A path π is *just* whenever, for every suffix of π beginning with state s and for every transition t enabled in state s , some transition u occurs in that suffix such that $t \not\smile u$. Equivalently, one might say that no enabled transition is denied forever only by concurrent transitions. The corresponding fairness assumption, which only allows just paths, is called *justness* (J).

Example 2 illustrates that J is strictly weaker than SC, i.e., J rules out fewer paths. While this system terminates under SC, it does not necessarily terminate under J, for it allows infinite communication between the *seller* and *buyer1*. Although the transition involving *buyer2* is relentlessly enabled, it is not ruled out by justness since the *seller* is involved in both communications.

Justness is however enough to assume that in our leading example at the top of the introduction, the two concurrent interactions cannot prevent each other from occurring.

Example 3: More formally, we can model the scenario described at the top of the introduction as follows.

$$\begin{aligned} & \text{seller1} \llbracket \mu X. \text{buyer1?order}; X \rrbracket \\ & \parallel \text{buyer1} \llbracket \mu Y. \text{seller1!order}; Y \rrbracket \\ & \parallel \text{seller2} \llbracket \mu Z. \text{buyer2?order}; Z \rrbracket \\ & \parallel \text{buyer2} \llbracket \mu W. \text{seller2!order}; W \rrbracket \end{aligned}$$

There is no just path where *seller2* and *buyer2* never act. Indeed, for any just path all components act infinitely often.

In general, $J \preceq WC$ holds [2]. In addition, for our session calculus, justness coincides with weak fairness of components.

Proposition 1: WC coincides with J.

Proof: Let π be an infinite path in our network that is not WC-fair. So, on a suffix of π , a component p is perpetually enabled, but never taken. In case p is stuck in a state where its next transition is a τ , then π is not just.

In case p is stuck in a state $q! \lambda; T$, then, for component p to be perpetually enabled, q must always be in a state $\sum_{i \in I} p_i? \lambda_i; T_i$ with $p = p_k$ and $\lambda = \lambda_k$ for some $k \in I$.

deadlock-freedom




Fig. 2. A classification for our session calculus of fairness assumptions and liveness properties.

Location q must get stuck in such a state, for if q keeps moving, it will at some point reach a state $\mu X.U$, which is not of the above form. Consequently, π is not just.

The remaining case is that p is stuck in a state of the form $\sum_{i \in I} p_i? \lambda_i; T_i$. For component p to be enabled, a component p_k with $k \in I$ must be in a state $p! \lambda_k; T$. Again it follows that π is not just. \square

As we will observe later, under a different choice of semantics of our session calculus, J and WC do not coincide.

B.3 Further notions of fairness: If we define *weak fairness of transitions* (WT), where, as for ST, the tasks are the individual transitions, then WT imposes no restrictions on the completed traces for our languages. To see why, observe that in any infinite path, no transition is enabled perpetually due to the τ -transitions for unfolding recursion. This most liberal fairness assumption, which we denote P,¹ only guarantees that the system as a whole will progress if some transition is enabled.

The survey [2] classifies 21 different notions of fairness, covering all common notions found in the literature. In our session calculus, many of these notions coincide, so that only 7 different notions of fairness remain; see Appendix A.

Here we have presented those that we found to be the most important notions for session calculi – summarised in Figure 2. Notably, there are strong fairness assumptions strictly between SC and ST. However, every fairness assumption from [2] leads to a notion of lock-freedom that coincides with one based on a fairness assumption defined in this section (see Appendix B).

C. A scheme for lock-freedom

As discussed, a fairness assumption rules out certain paths for given systems. As lock-freedom considers only the paths of a system that can actually be taken, it depends on the

¹*On terminology.* In related work [2], P stands for ‘‘progress’’, the assumption that a system cannot spontaneously halt as long as it is neither deadlocked nor successfully terminated. However, the word ‘‘progress’’ is heavily overloaded, meaning anything from deadlock-freedom [11], [12] and lock-freedom [5], [13], [14] to other liveness properties [15], such as weak and strong normalisation. That means, it refers to some desired property rather than an assumption on paths. Furthermore, there are related liveness properties such as *global progress* that concerns delegation [16].

underlying fairness assumption. Hence, a *scheme for lock-freedom* reads as follows:

Along any \mathcal{F} -fair path, if a component has not successfully terminated, then it must eventually do something. (1)

We can now formally define our scheme for lock-freedom with respect to a fairness assumption \mathcal{F} .

Definition 4: Let \mathcal{F} be a fairness assumption. A network \mathbb{N} satisfies liveness property $\mathcal{L}(\mathcal{F})$ (for short $\mathbb{N} \models \mathcal{L}(\mathcal{F})$) if, for each \mathcal{F} -fair path π starting in \mathbb{N} and each location p of \mathbb{N} ,

- either p successfully terminates on π , or
- π contains infinitely many transitions involving p .

Remember that a location p successfully terminates when it is of the form $p \llbracket \text{OK} \rrbracket$. The letter \mathcal{L} indicates “liveness” or “lock-freedom”.

We say $\mathcal{L}(\mathcal{F})$ is *stronger* than $\mathcal{L}(\mathcal{G})$, denoted by $\mathcal{L}(\mathcal{F}) \Rightarrow \mathcal{L}(\mathcal{G})$, if $\mathbb{N} \models \mathcal{L}(\mathcal{F}) \Rightarrow \mathbb{N} \models \mathcal{L}(\mathcal{G})$, for all \mathbb{N} . It is *strictly stronger* if moreover $\mathcal{L}(\mathcal{G}) \not\Rightarrow \mathcal{L}(\mathcal{F})$. In case a fairness assumption \mathcal{G} is stronger than \mathcal{F} , then $\mathcal{L}(\mathcal{G})$ is a weaker property than $\mathcal{L}(\mathcal{F})$.

Proposition 2: $\mathcal{F} \preceq \mathcal{G}$ implies $\mathcal{L}(\mathcal{F}) \Rightarrow \mathcal{L}(\mathcal{G})$, for fairness assumptions \mathcal{F} and \mathcal{G} .

Intuitively, any path of \mathbb{N} that is lock-free under \mathcal{F} will also be lock-free under \mathcal{G} . Since \mathcal{G} rules out more paths than \mathcal{F} and since \mathcal{L} is defined over paths, the proof is obvious.

A network has a *deadlock* (state) if there exists a reachable network state without outgoing transitions that is not successfully terminated; a network is *deadlock-free* if it does not have a deadlock.

Clearly, $\mathcal{L}(\text{ST})$ implies deadlock-freedom, since every finite execution can be extended to some ST-fair path, using feasibility. In networks consisting of one or two parties only, deadlock-freedom coincides with all notions of lock-freedom. Deadlock-freedom, however, is considered to be insufficient for networks with three or more locations, as those networks may experience starvation: *starvation* occurs when there is an execution path along which some component wants to perform a task but no task involving that component occurs.

Using the relationship between $\mathcal{L}(\text{ST})$ and deadlock-freedom, as well as Proposition 2, yields the classification of liveness properties on the right-hand side of Figure 2. Since $\mathcal{F} \not\preceq \mathcal{G}$ does not imply that $\mathcal{L}(\mathcal{F})$ is strictly stronger than $\mathcal{L}(\mathcal{G})$, we provide separating examples to prove that the presented notions of lock-freedom are different.

C.1 $\mathcal{L}(\text{ST})$ is strictly stronger than deadlock-freedom:

Example 4: Consider the following network, where a *buyer* purchases goods repeatedly from a *seller*, while a *shipper* is awaiting an order that is never placed.

$$\begin{aligned} & \text{buyer} \llbracket \mu X. \text{seller!buy}; X \rrbracket \\ \parallel & \text{seller} \llbracket \mu Y. \text{buyer?buy}; Y \rrbracket \\ \parallel & \text{shipper} \llbracket \text{seller?order} \rrbracket \end{aligned}$$

This network is deadlock-free, for the buyer and seller can always interact; it does not satisfy $\mathcal{L}(\text{ST})$ as *shipper* is not in state OK and is never involved in a transition.

C.2 $\mathcal{L}(\text{SC})$ is strictly stronger than $\mathcal{L}(\text{ST})$: Consider Example 1. We have seen that all ST-fair paths successfully terminate. In particular, along all fair paths the *shipper* performs a transition. In contrast, there is an infinite SC-fair path where the *shipper* neither makes a transition nor successfully terminates.

The following example separates ST from SC without considering termination.

Example 5: Consider the following network, where a *buyer* talks alternately to two *sellers*, but talks to each seller for as long as they desire.

$$\begin{aligned} & \text{buyer} \llbracket \mu X. (\text{seller1?talk}; X \\ & \quad + \text{seller1?wait}; \\ & \quad \mu Z. (\text{seller2?talk}; Z \\ & \quad \quad + \text{seller2?wait}; X)) \rrbracket \\ \parallel & \text{seller1} \llbracket \mu V. (\text{buyer!talk}; V \oplus \text{buyer!wait}; V) \rrbracket \\ \parallel & \text{seller2} \llbracket \mu W. (\text{buyer!talk}; W \oplus \text{buyer!wait}; W) \rrbracket \end{aligned}$$

The above network satisfies $\mathcal{L}(\text{ST})$ but not $\mathcal{L}(\text{SC})$, since no location terminates and there are SC-fair paths on which one of *seller1* or *seller2* ceases to act, violating the condition that there must be infinitely many transitions stemming from them.

C.3 $\mathcal{L}(\text{J})$ is strictly stronger than $\mathcal{L}(\text{SC})$: We consider a variant of Example 2.

Example 6:

$$\begin{aligned} & \text{seller} \llbracket \mu X. (\text{buyer1?order1}; X + \text{buyer2?order2}; X) \rrbracket \\ \parallel & \text{buyer1} \llbracket \mu Y. \text{seller!order1}; Y \rrbracket \\ \parallel & \text{buyer2} \llbracket \mu Z. \text{seller!order2}; Z \rrbracket \end{aligned}$$

The above network satisfies $\mathcal{L}(\text{SC})$, since each location p has a relentlessly enabled communication transition. Hence, p will engage in a communication transition infinitely often. However, the system does not satisfy $\mathcal{L}(\text{J})$, since there is a J-path where *buyer1* never acts. Namely, every communication of *buyer1* may be preempted by a communication of *buyer2*, as both buyers communicate with the same *seller*.

Although Example 2 separates J from SC, we cannot use it as separating example for $\mathcal{L}(\text{J})$ and $\mathcal{L}(\text{SC})$. It does not even satisfy $\mathcal{L}(\text{ST})$, since if *buyer2* ever acts, then *buyer1* never successfully terminates nor engages in infinitely many transitions.

C.4 $\mathcal{L}(\text{P})$ is strictly stronger than $\mathcal{L}(\text{J})$: The network of Example 3 – the example from the introduction – satisfies $\mathcal{L}(\text{J})$, since on a just path there are infinitely many transitions stemming from each location. However, it does not satisfy $\mathcal{L}(\text{P})$, since there exists a path where two components talk forever, to the exclusion of the other two. This example indicates (again) that J is the minimal realistic fairness assumption.

In Appendix B, we analyse further notions of lock-freedom, based on other fairness assumptions.

D. Lock-freedom in the literature

There are two prevalent notions of lock-freedom in the literature, which we call Kobayashi lock-freedom and Padovani

$$\begin{array}{c}
\frac{\mathbb{N} \equiv \mathbb{N}' \quad \mathbb{N}' \xrightarrow{\alpha} \mathbb{M}' \quad \mathbb{M}' \equiv \mathbb{M}}{\mathbb{N} \xrightarrow{\alpha} \mathbb{M}} \qquad \frac{p \llbracket \mathbb{T} \{ \mu X. \mathbb{T} / X \} \rrbracket \parallel \mathbb{N} \xrightarrow{\alpha} p \llbracket \mathbb{U} \rrbracket \parallel \mathbb{N}}{p \llbracket \mu X. \mathbb{T} \rrbracket \parallel \mathbb{N} \xrightarrow{\alpha} p \llbracket \mathbb{U} \rrbracket \parallel \mathbb{N}} \\
\\
\frac{j \in H \quad k \in I \quad \lambda_k = \lambda_j}{p_k \llbracket \bigoplus_{h \in H} q_h ! \lambda_h ; \mathbb{U}_h \rrbracket \parallel q_j \llbracket \sum_{i \in I} p_i ? \lambda_i ; \mathbb{T}_i \rrbracket \parallel \mathbb{N} \xrightarrow{p_k \rightarrow q_j : \lambda_k} p_k \llbracket \mathbb{U}_j \rrbracket \parallel q_j \llbracket \mathbb{T}_k \rrbracket \parallel \mathbb{N}}
\end{array}$$

Fig. 3. A reactive semantics without τ -transitions for internal choice or recursion. The definition of \equiv is unchanged.

lock-freedom, acknowledging the authors of key papers where these properties are investigated. We prove that these two notions relate to $\mathcal{L}(\text{SC})$ and $\mathcal{L}(\text{ST})$, respectively. We believe, however, that $\mathcal{L}(\text{J})$ is a novel notion of lock-freedom. In Section V we discuss further notions.

D.1 Kobayashi lock-freedom: Our scheme (1) for lock-freedom is inspired by a scheme proposed by Kobayashi [6] in the setting of the linear π -calculus, which does not feature operators for choice. Our scheme is more general, making it applicable to several calculi.

Although Kobayashi argues that lock-freedom is parametrised by a fairness assumption, he settles for exactly one, called *strong fairness* and attributed to [17], [18], with the stated intention that: “every process that is able to participate in a communication infinitely often can eventually participate in a communication.” The intended fairness assumption in [6] coincides with SC. Almost the same can be said for the formalisation of strong fairness in [6], although literally speaking the latter is slightly weaker.²

D.2 Padovani lock-freedom coincides with $\mathcal{L}(\text{ST})$: Padovani [5] presents a notion of lock-freedom that does not refer explicitly to a fairness assumption. Below we use the abbreviation $\text{PROC}(p, \mathbb{N})$ that denotes the unique thread state \mathbb{T} such that $\mathbb{N} \equiv p \llbracket \mathbb{T} \rrbracket \parallel \mathbb{N}'$, if p is a location of a network state \mathbb{N} .

Definition 5: \mathbb{N} is *Padovani lock-free* if for each reachable state \mathbb{M} of \mathbb{N} , and for each location p of \mathbb{M} such that $\text{PROC}(p, \mathbb{M}) \neq \text{OK}$, network \mathbb{M} has an execution path that contains a transition involving p .

Theorem 1: A network is Padovani lock-free iff it satisfies $\mathcal{L}(\text{ST})$. [See Appendix C for the proof.]

E. Lock-freedom for a reactive semantics

This section demonstrates that differences between session calculi, which may appear to be merely stylistic, in fact impact the resulting notions of liveness. An alternative semantics, (e.g. [7], [8]), which we call *reactive semantics*, is given in Figure 3. Here, neither unfolding recursion nor making a choice between various send actions induces a τ -transition.

²The reason is that Kobayashi’s intended requirement that a component must act is formalised by describing the states right before and right after that component acts, and stipulating that one must go from the former to the latter. However, in [6] there is no way to unambiguously project global states on individual components, and one can make the prescribed transition without actually involving that component.

In Definition 4, we formally introduced liveness properties for a network, parametrised by a fairness assumption. In fact, the definition also depends on the given semantics. In the remainder, we denote by $\mathcal{L}(\mathcal{F})$ a liveness property with regard to the semantics of Figure 1, and by $\mathcal{R}(\mathcal{F})$ a liveness property with regard to the reactive semantics.

Example 7: The following network has a deadlock by the default semantics of Figure 1. Consequently, it satisfies none of the properties $\mathcal{L}(\mathcal{F})$. Yet, it satisfies $\mathcal{R}(\mathcal{F})$, for each \mathcal{F} .

$$\begin{array}{l}
\text{buyer} \llbracket \text{seller!buy} \oplus \text{seller!order} \rrbracket \\
\parallel \text{seller} \llbracket \text{buyer?buy} \rrbracket
\end{array}$$

A similar result to Proposition 2 shows that the strength of a fairness assumption partially determines the strength of the corresponding liveness property.

Proposition 3: $\mathcal{F} \preceq \mathcal{G}$ implies $\mathcal{R}(\mathcal{F}) \Rightarrow \mathcal{R}(\mathcal{G})$, for fairness assumptions \mathcal{F} and \mathcal{G} .

Consequently, a classification of the liveness properties $\mathcal{R}(\mathcal{F})$, for \mathcal{F} any of the fairness assumptions from [2], can be obtained from the classification of these fairness properties (Figure 6 in Appendix A) by collapsing certain entries, just as for the classification of liveness properties $\mathcal{L}(\mathcal{F})$ from Figure 2. Since the separating examples given for $\mathcal{L}(\mathcal{F})$ apply also to $\mathcal{R}(\mathcal{F})$, we end up with at least four different notions $\mathcal{R}(\mathcal{F})$. However, we expect a lattice that is quite a bit larger, with fewer notions coinciding.

As an instance of this, $\mathcal{R}(\text{J})$ is strictly stronger than $\mathcal{R}(\text{WC})$. Strictness is shown by the following example.

Example 8: The following network presents a *buyer* who negotiates with *seller1* up to a point and then decides to order a product with *seller2* and inform *seller1* about their decision.

$$\begin{array}{l}
\text{buyer} \llbracket \mu X. (\text{seller1!negotiate}; X \\
\oplus \text{seller2!order}; \text{seller1!done}) \rrbracket \\
\parallel \text{seller1} \llbracket \mu Y. (\text{buyer?negotiate}; Y + \text{buyer?done}) \rrbracket \\
\parallel \text{seller2} \llbracket \text{buyer?order} \rrbracket
\end{array}$$

The network successfully terminates under $\mathcal{R}(\text{WC})$, for a transition involving *seller2* is perpetually enabled, when appealing to Figure 3. It does not need to terminate under justness as the *buyer* is involved in all transitions.

Similar to Example 2, guaranteeing termination in this example seems wrong as the fairness assumption constrains the *buyer*’s ‘free will’. Therefore, the presented results suggest that J is a more realistic notion than WC.

III. SESSION TYPES AND COMPLETENESS

We now focus on session type systems. A suitably crafted session type system guarantees liveness properties for a network, if the network is well-typed. We devise a session type system that is complete for $\mathcal{L}(J)$, meaning that all lock-free networks can be typed.

A. Global session types, projections and type judgements

We build on a widely-adopted approach for multiparty session types. It first defines a global type, describing the interacting behaviour of all parties involved. In our syntax for global types, communications of the form $p \rightarrow q : \lambda$ describe the sending of a message labelled λ from location p to q , and \boxplus indicates a choice over a finite, non-empty index set I .

$$\mathcal{G} := \begin{array}{ll} \text{OK} & \text{(successful termination)} \\ | \quad \boxplus_{i \in I} p \rightarrow q_i : \lambda_i ; \mathcal{G}_i & \text{(choice of communication)} \\ | \quad X & \text{(recursion variable)} \\ | \quad \mu X. \mathcal{G} & \text{(recursion)} \end{array}$$

As for our session type calculus we exclude types of the form $\mu X. X$ or $\mu X. \mu Y. \mathcal{G}$ to enforce guarded recursion. Moreover, for $\boxplus_{i \in I} p \rightarrow q_i : \lambda_i ; \mathcal{G}_i$, we assume $p \neq q_i$ for all $i \in I$. That means locations cannot send messages to themselves. A global type is *closed* whenever it contains no free recursion variables. The fact that p is the same in every branch of a choice means there is a distinguished choice leader p , who makes that choice, but there may be different recipients, as in related work on flexible choices [7].

A global session type can be projected to a local view for each location. We call local types stemming from projections *projection types*. They are defined almost in the same way as threads of Section II: instead of the construct $\sum_{i \in I} p_i ? \lambda_i ; T_i$ they feature merely its unary case $p ? \lambda ; T$, as well as the *merge* operators $\prod_{i \in I} T_i$.

We define the set of *participants* of a global type \mathcal{G} recursively:

$$\begin{aligned} \text{parties}(\text{OK}) &= \text{parties}(X) = \emptyset \\ \text{parties}(\mu X. \mathcal{G}) &= \text{parties}(\mathcal{G}) \\ \text{parties}(\boxplus_{i \in I} p \rightarrow q_i : \lambda_i ; \mathcal{G}_i) &= \bigcup_{i \in I} \{p, q_i\} \cup \text{parties}(\mathcal{G}_i) \end{aligned}$$

Given a global session type \mathcal{G} and location p , we define the projection $\mathcal{G}|_p$ of \mathcal{G} on p as follows.

$$\begin{aligned} \text{OK}|_p &= \text{OK} & X|_p &= X \\ (\mu X. \mathcal{G})|_p &= \begin{cases} \text{OK} & \text{if } p \notin \text{parties}(\mathcal{G}) \\ & \text{and } \mu X. \mathcal{G} \text{ is closed} \\ \mu X. (\mathcal{G}|_p) & \text{otherwise} \end{cases} \\ (\boxplus_{i \in I} p \rightarrow q_i : \lambda_i ; \mathcal{G}_i)|_r &= \begin{cases} \bigoplus_{i \in I} (p \rightarrow q_i : \lambda_i ; \mathcal{G}_i)|_r & p = r \\ \prod_{i \in I} (p \rightarrow q_i : \lambda_i ; \mathcal{G}_i)|_r & p \neq r \end{cases} \\ (p \rightarrow q : \lambda ; \mathcal{G})|_r &= \begin{cases} q ! \lambda ; (\mathcal{G}|_r) & p = r \\ p ? \lambda ; (\mathcal{G}|_r) & q = r \\ \mathcal{G}|_r & r \notin \{p, q\} \end{cases} \end{aligned}$$

The merge operator is interpreted directly through the judgement relation \vdash between threads and projection types, coinductively defined in Figure 4. See [19] for a formal definition

$$\begin{array}{c} \frac{\text{T}\{\mu X. T/X\} \vdash U}{\mu X. T \vdash U} \quad \frac{\text{T} \vdash U\{\mu X. U/X\}}{\text{T} \vdash \mu X. U} \\ \hline \text{OK} \vdash \text{OK} \quad \frac{i \in I \quad \text{T}_i \vdash U_i}{\sum_{i \in I} p_i ? \lambda_i ; \text{T}_i \vdash p_i ? \lambda_i ; U_i} \\ \hline \frac{I \subseteq J \quad \forall i \in I \quad \text{T}_i \vdash U_i}{\bigoplus_{i \in I} p_i ! \lambda_i ; \text{T}_i \vdash \bigoplus_{i \in J} p_i ! \lambda_i ; U_i} \quad \frac{\forall i \in I \quad \text{T} \vdash U_i}{\text{T} \vdash \prod_{i \in I} U_i} \end{array}$$

Fig. 4. Typing judgements relating threads to projection types.

of what it means to interpret such rules coinductively. Usually, the merge is defined independently from the type judgements; it is simply an operation that builds a single type from several types, without using an explicit merge primitive. In the standard approach [20], the work of our judgement relation \vdash is split between (a) the aforementioned merge operation, (b) a subtyping relation \leq between types [20, Definition 6], and (c) a relation \vdash between threads and local session types [20, Figure 5]. Our use of merge as a primitive construct for generating projection types, interpreted through \vdash , makes merging as general as possible.

B. Well-typed networks

The following definition plays the role of a type rule assigning a global type to a network in related systems, e.g., [7], [8], [20], [21].

Definition 6: A network $\mathbb{N} = p_1 \llbracket T_1 \rrbracket \parallel p_2 \llbracket T_2 \rrbracket \parallel \dots \parallel p_n \llbracket T_n \rrbracket$ is well-typed with respect to a global type \mathcal{G} , denoted $\mathbb{N} \vdash \mathcal{G}$, if \mathcal{G} is closed, $\text{parties}(\mathcal{G}) \subseteq \{p_1, p_2, \dots, p_n\}$, and $T_i \vdash \mathcal{G}|_{p_i}$ for all i .

A network \mathbb{N} is *well-typed* if $\mathbb{N} \vdash \mathcal{G}$ for some global type \mathcal{G} .

Example 9: A global type for the network of Example 1 is

$$\mathcal{G} = \mu X. (\text{buyer} \rightarrow \text{seller} : \text{talk} ; X \boxplus \text{buyer} \rightarrow \text{seller} : \text{buy} ; \text{seller} \rightarrow \text{shipper} : \text{order} ; \text{OK}).$$

We have

$$\begin{aligned} \mathcal{G}|_{\text{buyer}} &= \mu X. (\text{seller} ! \text{talk} ; X \oplus \text{seller} ! \text{buy} ; \text{OK}) \\ \mathcal{G}|_{\text{seller}} &= \mu X. (\text{buyer} ? \text{talk} ; X \sqcap \text{buyer} ? \text{buy} ; \text{shipper} ! \text{order} ; \text{OK}) \\ \mathcal{G}|_{\text{shipper}} &= \mu X. (X \sqcap \text{seller} ? \text{order} ; \text{OK}). \end{aligned}$$

With the help of the rules of Figure 4, we can derive the following facts, using proofs that are not well-founded.

$$\begin{aligned} \mu X. (\text{seller} ! \text{talk} ; X \oplus \text{seller} ! \text{buy} ; \text{OK}) \vdash \mathcal{G}|_{\text{buyer}} \\ \mu Y. (\text{buyer} ? \text{talk} ; Y \\ + \text{buyer} ? \text{buy} ; \text{shipper} ! \text{order} ; \text{OK}) \vdash \mathcal{G}|_{\text{seller}} \\ \text{seller} ? \text{order} ; \text{OK} \vdash \mathcal{G}|_{\text{shipper}} \end{aligned}$$

This network is well-typed. However, in the literature, it is commonly regarded as not well-typed, which may be due to the unguarded recursion in $\mathcal{G}|_{\text{shipper}}$

Example 10: This network is a restriction of the previous example, where no message is sent to the *shipper* on any path.

$$\begin{aligned} & \text{buyer} \llbracket \mu X.(\text{seller!talk}; X) \rrbracket \\ \parallel & \text{seller} \llbracket \mu Y.(\text{buyer?talk}; Y \\ & \quad + \text{buyer?buy}; \text{shipper!order}; \text{OK}) \rrbracket \\ \parallel & \text{shipper} \llbracket \text{seller?order}; \text{OK} \rrbracket \end{aligned}$$

By using the same global type \mathcal{G} , we can type the network. The following judgement makes use of the rule for internal choice in Figure 4, which permits deleting branches, as for most session subtype relations in the literature [22], [23].

$$\mu X.\text{seller!talk}; X \vdash \mathcal{G} \upharpoonright_{\text{buyer}}$$

The projections to the other locations are the same as in Example 9.

This network is well-typed and deadlock-free, but is not lock-free under any fairness assumption. That is, this type system is unsound for any notion of lock-freedom we have discussed.

Any type system targeting some notion of lock-freedom presented must reject Example 10. In this paper, our design decision to ensure soundness is to require that recursion has to be guarded for projections.³ We thereby disallow the projection $\mathcal{G} \upharpoonright_{\text{shipper}}$ in the above examples, thereby rejecting the networks in Examples 9 and 10. Since Example 9 satisfies $\mathcal{L}(\text{ST})$, we have to aim for a stronger notion of lock-freedom. It will be lock-freedom under justness and we will prove that our session type system is complete for that type of lock-freedom.

C. Guarded type judgements

We define a variant of well-typedness (Definition 6) that enforces each projection type to be guarded. A projection type T is *guarded* iff each occurrence of a variable X within a subexpression $\mu X.U$ of T occurs within a subexpression $p!\lambda; T$ or $p?\lambda; T$.

Definition 7: A network \mathbb{N} is *guardedly well-typed* with respect to a global type \mathcal{G} , denoted $\mathbb{N} \vdash^g \mathcal{G}$, if $\mathbb{N} \vdash \mathcal{G}$ and all projections $\mathcal{G} \upharpoonright_p$ are guarded.

Note that \mathcal{G} is guarded by definition, but this is not sufficient to ensure that $\mathcal{G} \upharpoonright_p$ is guarded. Examples 9 and 10 are well-typed, but not guardedly well-typed.

Example 11: Example 6, which features a competition between two buyers, is guardedly well-typed with respect to the following global type.

$$\mathcal{G} = \mu X.(\text{buyer1} \rightarrow \text{seller:order1}; \\ \text{buyer2} \rightarrow \text{seller:order2}; X)$$

All projections are guarded. Indeed, any global type without a choice will lead to guarded projections. The interesting

³An alternative design decision for strengthening the type system, that we do not pursue here, could be to restrict the type rule for internal choice (Figure 4) to prevent branches from being deleted (c.f. [7]), which, combined with our general merge, would allow Example 9 to stay in the fold for $\mathcal{L}(\text{ST})$.

projection relates the thread for the *seller* to the projection of the *seller*.

$$\begin{aligned} \mathcal{G} \upharpoonright_{\text{seller}} &= \mu X.(\text{buyer1?order1}; \text{buyer2?order2}; X) \\ &\mu X.(\text{buyer1?order1}; X + \text{buyer2?order2}; X) \vdash \mathcal{G} \upharpoonright_{\text{seller}} \end{aligned}$$

The above judgement holds by unfolding the recursions so as to appeal twice to the rule for \sum in Figure 4.

The following example illustrates that our type system cannot be complete for $\mathcal{L}(\text{SC})$.

Example 12: The next network satisfies $\mathcal{L}(\text{SC})$, but not $\mathcal{L}(\text{J})$.

$$\begin{aligned} & p \llbracket \mu X.(q!a; X \oplus q!b; X) \rrbracket \\ \parallel & q \llbracket \mu Y.(p?a; Y + r?c; (r?d; Y + p?b; r?d; Y)) \rrbracket \\ \parallel & r \llbracket \mu Z.q!c; q!d; Z \rrbracket \end{aligned}$$

The network does not satisfies $\mathcal{L}(\text{J})$ for there is an infinite just path in which locations p and q constantly communicate via $p \rightarrow q:a$ and r never engages in a communication. That just path is not SC-fair, since the communication $r \rightarrow q:c$ involving location r is relentlessly enabled yet never taken.

The network is not well-typed, let alone guardedly well-typed, for each global type must have a subexpression $p \rightarrow q:a; \mathcal{G}_1 \boxplus p \rightarrow q:b; \mathcal{G}_2$, and hence must have a reachable state \mathbb{M} in which both transitions $\mathbb{M} \xrightarrow{p \rightarrow q:a}$ and $\mathbb{M} \xrightarrow{p \rightarrow q:b}$ are enabled. Yet there is no such reachable state.

Example 12 shows that the strongest completeness result possible is completeness with respect to $\mathcal{L}(\text{J})$. Before turning to our completeness proof in the next section, we demonstrate the power of our general merge operator.

Example 13: This network consists of two independent pairs of threads, both of which make a choice repeatedly.

$$\begin{aligned} & \text{buyer1} \llbracket \mu X.(\text{seller1!wait}; X \oplus \text{seller1!order}) \rrbracket \\ \parallel & \text{seller1} \llbracket \mu Y.(\text{buyer1?wait}; Y + \text{buyer1?order}) \rrbracket \\ \parallel & \text{buyer2} \llbracket \mu X.(\text{seller2!wait}; X \oplus \text{seller2!order}) \rrbracket \\ \parallel & \text{seller2} \llbracket \mu Y.(\text{buyer2?wait}; Y + \text{buyer2?order}) \rrbracket \end{aligned}$$

The following is a global type for this example.

$$\mathcal{G} = \mu X. \left((\text{buyer1} \rightarrow \text{seller1:wait}; \\ (\text{buyer2} \rightarrow \text{seller2:wait}; X \\ \boxplus \text{buyer2} \rightarrow \text{seller2:order}; \mathcal{G}_Y)) \\ \boxplus \text{buyer1} \rightarrow \text{seller1:order}; \mathcal{G}_Z \right)$$

$$\begin{aligned} \text{with } \mathcal{G}_Y &= \mu Y.(\text{buyer1} \rightarrow \text{seller1:wait}; Y \\ &\quad \boxplus \text{buyer1} \rightarrow \text{seller1:order}; \text{OK}) \\ \mathcal{G}_Z &= \mu Z.(\text{buyer2} \rightarrow \text{seller2:wait}; Z \\ &\quad \boxplus \text{buyer2} \rightarrow \text{seller2:order}; \text{OK}). \end{aligned}$$

We can show that $\mathbb{N} \vdash^g \mathcal{G}$. For example, we have

$$\begin{aligned} \mathcal{G} \upharpoonright_{\text{buyer1}} &= \mu X.(\text{seller1!wait}; \\ &\quad (X \sqcap \mu Y.(\text{seller1!wait}; Y \\ &\quad \quad \boxplus \text{seller1!order}; \text{OK})) \\ &\quad \boxplus \text{seller1!order}; \text{OK}) \text{ and} \\ &\mu X.(\text{seller1!wait}; X \oplus \text{seller1!order}) \vdash \mathcal{G} \upharpoonright_{\text{buyer1}} \end{aligned}$$

where the projection $\mathcal{G} \upharpoonright_{\text{buyer1}}$ on *buyer1* is guarded.

$$\text{GT}(h, \mathbb{M}) = \begin{cases} \text{GT}(h, \mathbb{M}') & \text{if } \mathbb{M} \xrightarrow{\tau} \mathbb{M}' \text{ for a network } \mathbb{M}', \\ \text{OK} & \text{if } \text{PROC}(p, \mathbb{M}) = \text{OK} \text{ for each location } p \text{ of } \mathbb{M}, \\ \text{DEADLOCK} & \text{if no location is ready in } \mathbb{M}, \\ X_{\mathbb{M}} & \text{if } \mathbb{M} \text{ occurs in } h \text{ and } h \upharpoonright \mathbb{M} \text{ is complete for } \mathbb{M}, \\ \bigsqcup_{i \in I} p \rightarrow q_i : \lambda_i ; \text{GT}(h_i, \mathbb{M}_i^p) & \text{if } \mathbb{M} \text{ occurs in } h, p = \text{CH}(h, \mathbb{M}) \text{ and } \text{PROC}(p, \mathbb{M}) = \bigoplus_{i \in I} q_i ! \lambda_i ; T_i, \\ \mu X_{\mathbb{M}} . \bigsqcup_{i \in I} p \rightarrow q_i : \lambda_i ; \text{GT}(h_i, \mathbb{M}_i^p) & \text{if } p = \text{CH}(\varepsilon, \mathbb{M}) \text{ and } \text{PROC}(p, \mathbb{M}) = \bigoplus_{i \in I} q_i ! \lambda_i ; T_i. \end{cases}$$

where $h_i := h(\mathbb{M}, p, q_i)$, i.e., the sequence obtained from h by appending the triple (\mathbb{M}, p, q_i) .

Fig. 5. Algorithm for synthesising a global type for a network.

The above example is out of scope of most systems for global session types that do not feature an explicit parallel composition operator. These systems are incomplete, in the sense that there are lock-free networks that cannot be typed. Our session type system overcomes the incompleteness, due our general treatment of merge. A similar example, which also can be typed by our methodology, is given in [24]. It is used there to demonstrate that there are networks that cannot be typed using established notions of global type without parallel composition [25]. An alternative approach using coinductive projections has been proposed in [26].

D. Completeness for lock-freedom under justness

We now show one of our main results, namely that, for our session type calculus, all lock-free networks can be typed, when assuming justness. To the best of our knowledge, this is the first completeness result of this kind.

Theorem 2: If $\mathbb{N} \models \mathcal{L}(\mathbb{J})$, then \mathbb{N} is guardedly well-typed.

The proof (Appendix D) makes use of an algorithm for synthesising a global type from a network, along with a proof that the algorithm terminates with the correct guarded type.

To express the algorithm we require the following concepts. A *reachable network*, from a given network \mathbb{N} , is a reachable network state \mathbb{M} that happens to be a network, in the sense that $\text{PROC}(p, \mathbb{N}) \neq \top$ for all locations p of \mathbb{N} . A network \mathbb{M} is *unfolded* if there is no network \mathbb{M}' with $\mathbb{M} \xrightarrow{\tau} \mathbb{M}'$ (although there may be network states \mathbb{M}' with $\mathbb{M} \xrightarrow{\tau} \mathbb{M}'$). The unfolding of a network \mathbb{M} is the unique network \mathbb{M}' such that $\mathbb{M} (\xrightarrow{\tau})^* \mathbb{M}'$ and \mathbb{M}' is unfolded. A location p is *ready* in a network \mathbb{N} if $\text{PROC}(p, \mathbb{N}) = \bigoplus_{i \in I} q_i ! \lambda_i ; T_i$, and for each $i \in I$ there exists a transition $\mathbb{N} \xrightarrow{p \rightarrow q_i : \lambda_i} \mathbb{N}_i$, using the transition relation of Figure 3. Define a *history* as a sequence of triples (\mathbb{N}, p, q) with \mathbb{N} a network and p, q locations of \mathbb{N} . A history h is *complete* for a network \mathbb{M} if each location that is ready in \mathbb{M} occurs in h . If h is a history and \mathbb{M} a network expression that occurs in h , then $h \upharpoonright \mathbb{M}$ denotes the suffix of h that starts with the first occurrence of \mathbb{M} in h . Moreover, $h \upharpoonright \mathbb{M}$ denotes the prefix of h prior to the first occurrence of \mathbb{M} in h , so that $h = (h \upharpoonright \mathbb{M})(h \upharpoonright \mathbb{M})$. Call a location p *eligible* in a network state \mathbb{M} w.r.t. a history h if (a) p is ready in \mathbb{M} , and (b) either \mathbb{M} does not occur in h , or p does not occur in $h \upharpoonright \mathbb{M}$. Finally, DEADLOCK is a constant, temporarily added to the syntax of session types.

Our algorithm requires several choices. First, we select a fresh variable $X_{\mathbb{M}}$ for each unfolded network \mathbb{M} that is reachable from \mathbb{N} . We then pick a total order on the finite set of locations of \mathbb{N} , referred to as *age*, so that each nonempty set of locations has an oldest element. Finally, for each reachable network \mathbb{M} and each location p that is ready in \mathbb{M} , say with $\text{PROC}(p, \mathbb{M}) = \bigoplus_{i \in I} q_i ! \lambda_i ; T_i$, and for each $i \in I$, pick a network \mathbb{M}_i^p such that $\mathbb{M} \xrightarrow{p \rightarrow q_i : \lambda_i} \mathbb{M}_i^p$ and $\text{PROC}(p, \mathbb{M}_i^p) = T_i$.

Our algorithm employs the routine $\text{GT}(h, \mathbb{M})$, parametrised by the choice of a network \mathbb{M} and a history h , as defined in Figure 5. Here, $\text{CH}(h, \mathbb{M})$ is a partial function that selects, for a given history h and reachable network \mathbb{M} , the oldest location that is eligible in \mathbb{M} w.r.t. h . It is defined only when such a location exists. The case distinction in the figure is meant to be prioritised, in the sense that a later-listed option is taken only if none of the higher-listed options apply. Our algorithm is then defined to yield the global type $\text{GT}(\varepsilon, \mathbb{N})$, with \mathbb{N} the given network and ε the empty history (sequence).

The intuition for our algorithm, which attempts to construct a global session type \mathcal{G} out of a given network \mathbb{N} , is as follows. Since it is essential that \mathcal{G} induces ongoing progress of all untermiated locations in the network, we keep track of the history h of communications incorporated in \mathcal{G} until the “construction front” at network state \mathbb{M} . Here the routine $\text{GT}(h, \mathbb{M})$ specifies the next communication-choice that will be incorporated in \mathcal{G} . The first clause in Figure 5, where the τ -transition must unfold recursion, says that all recursions should be unfolded before attempting the remaining case distinctions. The second clause says that we can safely terminate upon reaching a state in which the threads of all locations have terminated. The last two clauses specify a choice leader p and extend \mathcal{G} with the send actions of p in state \mathbb{M} . Here p must be a location that is ready in \mathbb{M} ; if such a p does not exist the failed attempt is reported by including the constant DEADLOCK in the attempted session type \mathcal{G} (Clause 3). Since the syntactic expression \mathcal{G} must be finite, each branch that does not reach OK needs to loop back to a previous stage in the construction of \mathcal{G} , at some point. To facilitate looping back, we attach a recursion variable $X_{\mathbb{M}}$ to each stage we might want to loop back to, namely to each first occurrence of a network state \mathbb{M} in our history. This explains the difference between Clauses 5 and 6. Clause 4 says that we can safely loop back to a previous stage if it involves the same current network state \mathbb{M} , and between that previous stage and the present all

locations that are ready in \mathbb{M} already had a turn. If Clause 4 does not apply, then \mathbb{M} must have eligible locations w.r.t. h , and Clause 5 or 6 picks the oldest such location, to make sure that in the end all eligible locations get a turn.

Example 14: Applying this algorithm to the network of Example 1 yields the type of Example 9, but with a spurious recursive anchor μZ right before $seller \rightarrow shipper:order$. Applied to Example 2 it fails with possible output

$$\begin{aligned} &\mu X. buyer1 \rightarrow seller:order1; \\ &buyer2 \rightarrow seller:order2; DEADLOCK. \end{aligned}$$

For Example 3 it yields the following correct type.

$$\mu X. (buyer1 \rightarrow seller1:order; buyer2 \rightarrow seller2:order; X)$$

For Example 4 it yields the type $\mu X. buyer \rightarrow seller:buy; X$; this type is incorrect for that network \mathbb{N} . This does not contradict the proof of Theorem 2, since $\mathbb{N} \not\models \mathcal{L}(J)$.

For Example 6 the algorithm yields the type of Example 11. For Example 7 it fails with output DEADLOCK.

For Example 8 it yields the following correct type.

$$\begin{aligned} &\mu X. (buyer \rightarrow seller1:negotiate; X \\ &\boxplus buyer \rightarrow seller2:order; \mu Z. buyer \rightarrow seller1:done; OK) \end{aligned}$$

For Example 13 it yields the type of Example 13.

Observation 1: An immediate consequence of Proposition 1 and Theorem 2, is that guardedly well-typed networks are complete for $\mathcal{L}(WC)$, suggesting that a carefully selected notion of weak fairness is suitable for some session calculi.

Corollary 1: If $\mathbb{N} \models \mathcal{L}(WC)$ then \mathbb{N} is guardedly well-typed.

In contrast, recall that Example 7 satisfies $\mathcal{R}(P)$; yet it is not (guardedly) well-typed. This shows that there is no corresponding completeness result for any notion of lock-freedom $\mathcal{R}(\mathcal{F})$, where \mathcal{F} is some notion of fairness. This is an argument for why we emphasise the semantics in Figure 1 rather than the one in Figure 3.

IV. RACE-FREEDOM AND SOUNDNESS

We have established that completeness holds, with respect to $\mathcal{L}(J)$. Hence, if we can model-check $\mathcal{L}(J)$, we know we can always synthesise a global type for a network. In this section, we consider soundness, meaning that a network is lock-free if it is (guardedly) well-typed. To complement our completeness result, we target $\mathcal{L}(J)$ and prove soundness for guardedly well-typed networks that are additionally race-free. The insight of this section is that soundness can be achieved when making the minimal fairness assumption justness.

Definition 8: A network state \mathbb{N} has a *race* whenever $\mathbb{N} \xrightarrow{p \rightarrow r: \lambda} \mathbb{N}'$ and $\mathbb{N} \xrightarrow{q \rightarrow r: \mu} \mathbb{N}''$ with either $p \neq q$ or $\mathbb{N}' \neq \mathbb{N}''$. A network is *race-free* if it has no reachable network state with a race.

Figure 2 implies that $\mathcal{L}(J)$ is the strongest lock-freedom property we can get, with the exception of $\mathcal{L}(P)$. Our guarded type system cannot be sound for the latter notion of lock-freedom, for Example 3 is guardedly well-typed and race-free, but does not satisfy $\mathcal{L}(P)$.

Our example to distinguish J and SC features races. Indeed, there is no race-free network separating J from SC, as confirmed by the following proposition.

Proposition 4: On race-free networks, J coincides with SC, for our session calculus in Figure 1.

Proof: Let π be an infinite path in a network that is not SC-fair. So, on π , a component p is infinitely often enabled, but never taken.

In case p is stuck in a state where its next transition is a τ , then π is not just.

In case p is stuck in a state ${}^q!\lambda; P$, then, in the first state on π on which p is enabled, q must be in a state $\sum_{i \in I} p_i? \lambda_i; T_i$ with $p = p_k$ and $\lambda = \lambda_k$ for some $k \in I$. If q remains in this state throughout π , then π is not just. If q leaves this state via a transition of π that does not involve p , then the state where this happens must be a race, and the network is not race-free.

In the remaining case, p is stuck in a state $\sum_{i \in I} p_i? \lambda_i; T_i$. For p to be enabled, a component p_k with $k \in I$ must be in a state $p! \lambda_k; P$, which reduces this case to the previous one. \square

Observation 2: In contrast to Proposition 4, for the session calculus in Figure 3, J and SC do not coincide, even for race-free networks (race-freedom also needs to be reformulated for that semantics). Example 8 illustrates this fact.

Since Examples 1, 3 and 4 are race-free, the remaining four notions $\mathcal{L}(P)$, $\mathcal{L}(SC)$, $\mathcal{L}(ST)$ and deadlock-freedom from Figure 2 are all different for race-free networks. Consequently, for race-free networks, the only collapse of Figure 2 is $\mathcal{L}(J) \Leftrightarrow \mathcal{L}(SC)$.

Example 15: The following network is race-free, guardedly well-typed and satisfies $\mathcal{L}(J)$. It is a variant of Example 8, which is also race-free, but does not satisfy $\mathcal{L}(J)$.

$$\begin{aligned} &buyer \llbracket \mu X. ((seller1!order1; seller2!wait; X) \\ &\quad \oplus (seller2!order2; seller1!done)) \rrbracket \\ &\parallel seller1 \llbracket \mu Y. (buyer?order1; Y + buyer?done) \rrbracket \\ &\parallel seller2 \llbracket \mu Z. (buyer?wait; Z + buyer?order2) \rrbracket \end{aligned}$$

This network is well-typed, for we can use the following global type.

$$\begin{aligned} \mathcal{G} = &\mu X. ((buyer \rightarrow seller1:order1; \\ &buyer \rightarrow seller2:wait; X) \\ &\boxplus (buyer \rightarrow seller2:order2; \\ &buyer \rightarrow seller1:done)) \end{aligned}$$

This example illustrates that it is possible to send messages to several locations via an internal choice in a race-free way. A simple way to prevent races, adopted by many session type systems, e.g. [20], [21], [27], is to ensure that each location listens to only one other location at any time. To achieve this we can impose the following syntactic restriction.

Definition 9: A network \mathbb{N} is *syntactically race-free*, if, for every sub-expression of the form $\sum_{i \in I} p_i? \lambda_i; T_i$, we have $p_i = p_j$ for all $i, j \in I$, and $\lambda_i \neq \lambda_j$ for all $i, j \in I$ with $i \neq j$.

Restricting external choices this way was never intended to be a complete criteria for race-freedom. However, it is a cheap

linear syntactic property to check, whereas checking for race-freedom is as complex as checking for deadlock-freedom.

Example 16: The following example is race-free, but not syntactically race-free.

$$\begin{aligned} & \text{buyer} \llbracket \mu X. (\text{seller}! \text{order}1; X \oplus \text{seller}2! \text{order}2) \rrbracket \\ \parallel & \text{seller}1 \llbracket \mu Y. (\text{buyer}? \text{order}1; \text{seller}2! \text{wait}; Y \\ & \quad + \text{seller}2? \text{done}) \rrbracket \\ \parallel & \text{seller}2 \llbracket \mu Z. (\text{seller}1? \text{wait}; Z \\ & \quad + \text{buyer}? \text{order}2; \text{seller}1! \text{done}) \rrbracket \end{aligned}$$

It is also guardedly well-typed and satisfies $\mathcal{L}(\mathbb{J})$. Soundness results are stronger if race-free networks are considered, rather than syntactically race-free networks.

The reverse of Theorem 2 does not hold. Hence we cannot expect a soundness result for all networks. It is not even the case that guardedly well-typed networks are deadlock-free.

Example 17: The following network is guardedly well-typed, but neither race-free nor deadlock-free, and hence certainly not $\mathcal{L}(\mathbb{J})$.

$$\begin{aligned} & \text{buyer}1 \llbracket \text{seller}! \text{buy}1; \text{OK} \rrbracket \\ \parallel & \text{buyer}2 \llbracket \text{seller}! \text{buy}2; \text{OK} \rrbracket \\ \parallel & \text{seller} \llbracket (\text{buyer}1? \text{buy}1; \text{buyer}2? \text{buy}2; \text{OK}) \\ & \quad + (\text{buyer}2? \text{buy}2; \text{buyer}1? \text{buy}1; \\ & \quad \text{buyer}1! \text{order}; \text{OK}) \rrbracket \end{aligned}$$

The global type for this example is the following.

$$\mathcal{G} = \text{buyer}1 \rightarrow \text{seller} : \text{buy}1 ; \text{buyer}2 \rightarrow \text{seller} : \text{buy}2 ; \text{OK}$$

In particular, $\mathcal{G} \upharpoonright_{\text{seller}} = \text{buyer}1? \text{buy}1 ; \text{buyer}2? \text{buy}2$ and

$$\begin{aligned} & \text{buyer}1? \text{buy}1 ; \text{buyer}2? \text{buy}2 \\ & + \text{buyer}2? \text{buy}2 ; \text{buyer}1? \text{buy}1 ; \text{buyer}1! \text{order} \vdash^g \mathcal{G} \upharpoonright_{\text{seller}} \end{aligned}$$

which holds due to the rule for \sum in Figure 4, permitting branches of an external choice to be removed.

This example may suggest that the culprit preventing soundness is the flexible external choice, i.e., the subtype relation \vdash^g . However, even if \vdash^g would be almost the identity relation, with each merge on types corresponding to an external choice of the corresponding threads, there would be guardedly well-typed networks that are not deadlock-free.

Example 18: Consider the following network.

$$\begin{aligned} & p \llbracket (s!a; t!a; r!d) \oplus (s!b; t!b) \rrbracket \\ \parallel & r \llbracket (s?c; t?e; p?d) + (t?e; s?c) \rrbracket \\ \parallel & s \llbracket p?a; r!c + p?b; r!c \rrbracket \\ \parallel & t \llbracket p?a; r!e + p?b; r!e \rrbracket \end{aligned}$$

Using the global type

$$\begin{aligned} \mathcal{G} = & (p \rightarrow s : a ; p \rightarrow t : a ; s \rightarrow r : c ; t \rightarrow r : e ; p \rightarrow r : d ; \text{OK}) \\ & \boxplus p \rightarrow s : b ; p \rightarrow t : b ; t \rightarrow r : e ; s \rightarrow r : c ; \text{OK} \end{aligned}$$

yields projections that are identical to the network threads, e.g.

$$\mathcal{G} \upharpoonright_p = (s!a; t!a; r!d) \oplus s!b; t!b.$$

So, $\mathbb{N} \vdash^g \mathcal{G}$ follows by using the identity as subtyping relation. Yet, this network is not deadlock-free, for the execution $p \rightarrow s : b ; s \rightarrow r : c ; p \rightarrow t : b ; t \rightarrow r : e$ reaches a deadlock with hanging input $p?d$ in location r .⁴

Examples 17 and 18 are both excluded by race-freedom. We now prove another main result, namely that the converse of Theorem 2 holds for race-free networks.

Theorem 3: If \mathbb{N} is guardedly well-typed and race-free, then $\mathbb{N} \models \mathcal{L}(\mathbb{J})$.

The proof [see Appendix E] hinges on the following *session fidelity* result, for which we appeal to race-freedom:

For race-free network states \mathbb{N} , if $\mathbb{N} \xrightarrow{p \rightarrow q : \lambda} \mathbb{M}$ and $\mathbb{N} \vdash^g \mathcal{G}$ then there exists \mathcal{G}' such that $\mathcal{G} \xrightarrow{p \rightarrow q : \lambda} \mathcal{G}'$ and $\mathbb{M} \vdash^g \mathcal{G}'$. Here, $\xrightarrow{\alpha}$ is a transition relation on global types, defined for this purpose. Session fidelity strengthens subject reduction, by insisting that the form of \mathcal{G}' reflects the transition.

From this statement we conclude that each reachable network state \mathbb{N}' along any just path π is guardedly well-typed. For every location p , either there are infinitely many transitions along π involving p , or there exists a suffix π' of π stemming from network state \mathbb{N}' , such that location p has no further transition involving p . Using justness and the fact that \mathbb{N}' is guardedly well-typed one can show that in the latter case p has successfully terminated.

Using Theorems 2 and 3, and Proposition 4 leads to a soundness and a completeness result for our type system.

Corollary 2: For race-free network \mathbb{N} , \mathbb{N} is guardedly well-typed iff \mathbb{N} satisfies $\mathcal{L}(\text{SC})$.

Observation 3: Projections are defined such that recursion maps to OK whenever $p \notin \text{parties}(\mathcal{G})$ and $\mu X. \mathcal{G}$ is closed (see Page 7). The latter condition plays an essential role for soundness. To see why, consider the following global type.

$$\mu X. p \rightarrow q : a ; \mu Y. r \rightarrow q : b ; X$$

The anchor with variable Y should, intuitively, be useless. However, if we were to exclude condition “ $\mu X. \mathcal{G}$ is closed”, the above global type would type the following network.

$$\begin{aligned} & p \llbracket \mu X. q!a ; \text{OK} \rrbracket \\ \parallel & q \llbracket \mu X. p?a ; \mu Y. r?b ; X \rrbracket \\ \parallel & r \llbracket \mu X. \mu Y. q!b ; X \rrbracket \end{aligned}$$

This race-free network reaches a deadlock right after communications $p \rightarrow q : a ; r \rightarrow q : b$. The above closedness-condition resolves this issue, and can be used to correct papers on global types featuring recursion binders.

V. RELATED AND FUTURE WORK ON LOCK-FREEDOM

While our completeness result is the first of its kind, there are several soundness results for type systems with respect to some notion of lock-freedom, e.g., [5], [8], [13], [29], [30], the most closely related of which we draw attention to in this

⁴ This is a counterexample to subject reduction in previous work that allows multiple recipients in an external choice [7], [28]. Those soundness results can be restored by restricting to race-free networks.

section. We also situate related work on lock-freedom with regard to our classification and point to future challenges.

a) Strong lock-freedom: Severi and Dezani-Ciancaglini propose a notion of *strong lock-freedom* [8] that coincides with $\mathcal{R}(J)$ for race-free networks. They employ a reactive semantics. The authors impose a restriction on paths, ensuring that all concurrent transitions proceed in lockstep. Their assumption is not a fairness assumption, as defined here, as it does not satisfy feasibility. However, it does have the effect of assuming justness, up to permutations of transitions, for race-free networks. By Observation 1, a completeness result along the lines of Theorem 2 cannot hold for strong lock-freedom. Strong lock-freedom cannot be lifted directly to networks with races. A minimal change to their definitions requiring a maximal number of enabled locations to act in every step would extend their definition to networks with races; we did not analyse this extension. Their use of coinductive syntax, rather than binders, is an alternative for avoiding the soundness problem in Observation 3 that is common in the literature.

b) Further lock-freedom schemes: Carbone, Dardha and Montesi translate Kobayashi’s scheme for lock-freedom to a session calculus where both internal and external choices are with a single location [30]. Their scheme is instantiated with SC^5 and coincides with $\mathcal{L}(SC)$, restricted to their calculus. Their scheme inherits the ambiguity discussed in Section II-D1. It assumes a semantics intermediate to those we study in this work, where internal choice is like in Figure 1, but recursion and singleton internal choice are reactive as in Figure 3. This makes their approach weaker than ours, if lifted directly to our calculus with flexible choices: Example 6 is lock-free under their scheme instantiated with assumption J (or even P), but does not satisfy $\mathcal{L}(J)$ (or even $\mathcal{R}(J)$) in our scheme. Note that their work concerns binary session types with delegation, which we do not consider.

Scalas and Yoshida propose the notions of $LIVE$, $LIVE+$, and $LIVE++$ [24]. The first, $LIVE$, follows the scheme of Padovani, hence coincides with $\mathcal{L}(ST)$. The second, $LIVE+$, is essentially another formulation of $\mathcal{L}(SC)$. The third, $LIVE++$, coincides with $\mathcal{L}(P)$, hence is unsound for session calculi since it rejects key examples such as Example 3. The definitions of [24] are arguably less portable than Definition 5, since their definitions refer to specific language features.

c) Asynchronous session calculi: An evaluation of lock-freedom for asynchronous calculi, where queues are inserted between communicating threads, requires separate attention. The asynchronous analogue to our session calculus is an infinite-state system. Therefore, ST is no longer the strongest fairness assumption; this is now *full fairness* (Fu) [2]. At the other end of the spectrum, there are also complications when defining concurrency of transitions (Definition 2). It is a design decision whether a thread is treated as a single component along with its queues, and whether enqueue and

dequeue events for the same queue are dependent or concurrent. Consequently, synchrony/asynchrony and the spectrum of fairness assumptions are not entirely perpendicular dimensions when defining notions of lock-freedom. Fairness plays an essential role in related work on preciseness of subtyping for asynchronous calculi [31], which is further evidence that fairness assumptions require scrutiny here.

d) Synthesis and multiparty compatibility: The body of literature on synthesising global types from multiparty compatible local types [25], [32], [33] plays a complementary role to our synthesis results, used to establish completeness. Usually, it is immediate that networks inhabiting a global type are multiparty compatible. Hence, we expect that a corollary of Theorem 2 is that $\mathcal{L}(J)$ implies multiparty compatibility, for some notion of multiparty compatibility. If we further assume a synthesis result showing that multiparty compatible networks are guardedly well-typed – under conditions such as race-freedom – then that notion of multiparty compatibility coincides with $\mathcal{L}(J)$. Such a result for our global type system, does not quite follow immediately from synthesis results in the literature, since Example 13 would require parallel composition in related work. The formal development of multiparty compatibility is left as future work.

e) Fair subtyping and weak normalisation: A fair subtyping relation has been defined for session types as the largest relation over threads that preserves weak normalisation [34]. Weak normalisation is the property that, at any point during an execution, it is not inevitable that a network will not successfully terminate. Weak normalisation is strictly stronger than Padovani’s notion of lock-freedom (Definition 5) – since the possibility of all components to successfully terminate entails the possibility of all components performing some enabled action – but is incomparable to liveness properties stronger than $\mathcal{L}(SC)$, including $\mathcal{L}(J)$ – Example 1 is weakly normalising, but does not satisfy $\mathcal{L}(SC)$. Consequently, the proposed notion of fair subtyping does not quite fit lock-freedom. Investigating a notion of fair subtyping that is adequate for $\mathcal{L}(ST)$ rather than weak normalisation, and also identifying a session type system complete for $\mathcal{L}(ST)$, as hinted at in the discussion surrounding Example 10, is future work. In particular, we do not claim that $\mathcal{L}(J)$ is the only notion of lock-freedom that can be characterised by some session type system.

VI. CONCLUSION

In this paper, we have systematically classified the notions of lock-freedom that arise by taking every fairness assumption listed in a recent survey [2]. Based on our comprehensive analysis, we are compelled to put forward a notion of fairness suitable for session calculi: *justness* (Definition 3), and its resulting notion of lock-freedom $\mathcal{L}(J)$, which we propose to call “*just lock-freedom*”. Through a generalisation of the classical merge operation on local session types, we have devised a session type system that is complete for just lock-freedom. Moreover, race-free networks are sound for just lock-

⁵The authors do not provide a definition of fairness, but cite Kobayashi [6] instead. Kobayashi’s definition does not lift immediately to the calculus of [30]. However, the authors appear to intend SC .

freedom. Justness is always reasonable to assume, since it does not constrain the ‘free will’ of participants (c.f. Examples 2 and 8), while ensuring that concurrent transitions do not constrain each other (c.f. Examples 3 and 13).

A strength of our results is that completeness (Theorem 2) holds for networks with flexible choice, in which branches of the same choice operator may involve different locations. Completeness suggests a methodology for session calculi that allows us to pass straight from any network satisfying our realistic notion of lock-freedom $\mathcal{L}(J)$ to a global session type. The methodology would be to directly model check that a network satisfies $\mathcal{L}(J)$, and then use the algorithm in Figure 5 to synthesise a global type for that network. This methodology works even for networks featuring races.

Interestingly, there are no previous results synthesising global types directly from lock-freedom. Indeed, Example 13, which satisfies almost all notions of lock-freedom in the literature, is known to be out of scope of related session type systems based on global types without explicit parallel composition. While this incompleteness issue in related work is partly due to the less general merge operator employed in those systems, another reason that enables us to obtain the completeness result in Theorem 2 is our scrutiny of the role of fairness assumptions. In fact, Example 12 shows that completeness of our session type system cannot be attained when assuming strong fairness of components. Furthermore, even small variations in the choice of semantics for the transition system can affect fairness assumptions significantly, weakening corresponding notions of lock-freedom (see Observation 1). Indeed, amongst all notions of lock-freedom considered in this paper, only $\mathcal{L}(J)$ yields both completeness for all networks and soundness for race-free networks (Theorem 3).

Acknowledgement: The key question addressed in this paper arose in conversation with Ilaria Castellani, Mariangiola Dezani-Ciancaglini, and Paola Giannini, to whom we are grateful for their generous feedback on this work.

REFERENCES

- [1] S. S. Owicki and L. Lamport, “Proving liveness properties of concurrent programs,” *ACM Trans. Program. Lang. Syst.*, vol. 4, no. 3, pp. 455–495, 1982. Available: <https://doi.org/10.1145/357172.357178>
- [2] R. J. van Glabbeek and P. Höfner, “Progress, justness, and fairness,” *ACM Computing Surveys*, vol. 52, no. 4, 2019. Available: <https://doi.org/10.1145/3329125>
- [3] S. D. Brookes, C. A. R. Hoare, and A. W. Roscoe, “A theory of communicating sequential processes,” *Journal of the ACM*, vol. 31, no. 3, pp. 560–599, 1984. Available: <https://doi.org/10.1145/828.833>
- [4] R. De Nicola and M. Hennessy, “CCS without τ ’s,” in *TAPSOFT ’87*, H. Ehrig, R. Kowalski, G. Levi, and U. Montanari, Eds. Springer, 1987, pp. 138–152. Available: https://doi.org/10.1007/3-540-17660-8_53
- [5] L. Padovani, “Deadlock and lock freedom in the linear π -calculus,” in *CSL-LICS ’14*, T. A. Henzinger and D. Miller, Eds. ACM, 2014. Available: <https://doi.org/10.1145/2603088.2603116>
- [6] N. Kobayashi, “A type system for lock-free processes,” *Inf. Comput.*, vol. 177, no. 2, pp. 122–159, 2002. Available: <https://doi.org/10.1006/inco.2002.3171>
- [7] I. Castellani, M. Dezani-Ciancaglini, and P. Giannini, “Reversible sessions with flexible choices,” *Acta Informatica*, vol. 56, no. 7-8, pp. 553–583, 2019. Available: <https://doi.org/10.1007/s00236-019-00332-y>
- [8] P. Severi and M. Dezani-Ciancaglini, “Observational equivalence for multiparty sessions,” *Fundam. Inform.*, vol. 170, no. 1-3, pp. 267–305, 2019. Available: <https://doi.org/10.3233/FI-2019-1863>

- [9] S. Jongmans and N. Yoshida, “Exploring type-level bisimilarity towards more expressive multiparty session types,” in *ESOP ’20*, P. Müller, Ed. Springer, 2020, pp. 251–279. Available: https://doi.org/10.1007/978-3-030-44914-8_10
- [10] K. R. Apt, N. Francez, and S. Katz, “Appraising fairness in languages for distributed programming,” *Distributed Computing*, vol. 2, pp. 226–241, 1988. Available: <https://doi.org/10.1007/BF01872848>
- [11] K. Honda, N. Yoshida, and M. Carbone, “Multiparty asynchronous session types,” *Journal of the ACM*, vol. 63, no. 1, pp. 9:1–9:67, 2016. Available: <https://doi.org/10.1145/2827695>
- [12] M. Dezani-Ciancaglini, D. Mostrous, N. Yoshida, and S. Drossopoulou, “Session types for object-oriented languages,” in *ECOOP ’06*, D. Thomas, Ed. Springer, 2006, pp. 328–352. Available: https://doi.org/10.1007/11785477_20
- [13] M. Dezani-Ciancaglini, N. Yoshida, A. J. Ahern, and S. Drossopoulou, “A distributed object-oriented language with session types,” in *Trustworthy Global Computing, International Symposium, TGC ’05, Revised Selected Papers*, R. De Nicola and D. Sangiorgi, Eds. Springer, 2005, pp. 299–318. Available: https://doi.org/10.1007/11580850_16
- [14] E. Najm, A. Nimour, and J. Stefani, “Guaranteeing liveness in an object calculus through behavioural typing,” in *FORTE XII / PSTV XIX*, J. Wu, S. T. Chanson, and Q. Gao, Eds. Kluwer, 1999, pp. 203–221. Available: https://doi.org/10.1007/978-0-387-35578-8_12
- [15] J. Misra, *A Discipline of Multiprogramming: Programming Theory for Distributed Applications*. Springer, 2001, ch. Progress Properties, pp. 155–213. Available: https://doi.org/10.1007/978-1-4419-8528-6_6
- [16] M. Coppo, M. Dezani-Ciancaglini, N. Yoshida, and L. Padovani, “Global progress for dynamically interleaved multiparty sessions,” *Mathematical Structures in Computer Science*, vol. 26, no. 2, pp. 238–302, 2016. Available: <https://doi.org/10.1017/S0960129514000188>
- [17] G. Costa and C. Stirling, “Weak and strong fairness in CCS,” *Information and Computation*, vol. 73, no. 3, pp. 207–244, 1987. Available: [https://doi.org/10.1016/0890-5401\(87\)90013-7](https://doi.org/10.1016/0890-5401(87)90013-7)
- [18] E. A. Emerson, “Temporal and modal logic,” in *Handbook of Theoretical Computer Science (vol. B): Formal Models and Semantics*. MIT press, 1990, pp. 995–1072. Available: <https://dl.acm.org/doi/10.5555/114891.114907>
- [19] R. J. van Glabbeek, “Coinductive validity.” Available: <http://arxiv.org/abs/2104.13021>
- [20] N. Yoshida and L. Gheri, “A very gentle introduction to multiparty session types,” in *Distributed Computing and Internet Technology*, D. V. Hung and M. D’Souza, Eds. Springer, 2020, pp. 73–93. Available: https://doi.org/10.1007/978-3-030-36987-3_5
- [21] P.-M. Denielou, N. Yoshida, A. Bejleri, and R. Hu, “Parameterised Multiparty Session Types,” *Log. Meth. Comp. Sci.*, vol. Volume 8, Issue 4, 2012. Available: [https://doi.org/10.2168/LMCS-8\(4:6\)2012](https://doi.org/10.2168/LMCS-8(4:6)2012)
- [22] S. J. Gay and M. Hole, “Subtyping for session types in the pi calculus,” *Acta Informatica*, vol. 42, no. 2, pp. 191–225, 2005. Available: <https://doi.org/10.1007/s00236-005-0177-z>
- [23] R. Demangeon and K. Honda, “Full abstraction in a subtyped pi-calculus with linear types,” in *CONCUR ’11*, J.-P. Katoen and B. König, Eds. Springer, 2011, pp. 280–296. Available: https://doi.org/10.1007/978-3-642-23217-6_19
- [24] A. Scalas and N. Yoshida, “Less is more: multiparty session types revisited,” *PACMPL*, vol. 3, no. POPL, pp. 30:1–30:29, 2019. Available: <https://doi.org/10.1145/3290343>
- [25] J. Lange, E. Tuosto, and N. Yoshida, “From communicating machines to graphical choreographies,” in *POPL ’15*. ACM, 2015, pp. 221–232. Available: <https://doi.org/10.1145/2676726.2676964>
- [26] F. Barbanera, M. Dezani-Ciancaglini, I. Lanese, and E. Tuosto, “Composition and decomposition of multiparty sessions,” *Journal of Logical and Algebraic Methods in Programming*, vol. 119, 100620, 2021. Available: <https://doi.org/10.1016/j.jlamp.2020.100620>
- [27] S. Ghilezan, S. Jakšić, J. Pantović, A. Scalas, and N. Yoshida, “Precise subtyping for synchronous multiparty sessions,” *Journal of Logical and Algebraic Methods in Programming*, vol. 104, pp. 127–173, 2019. Available: <https://doi.org/10.1016/j.jlamp.2018.12.002>
- [28] I. Castellani, M. Dezani-Ciancaglini, P. Giannini, and R. Horne, “Global types with internal delegation,” *Theoretical Computer Science*, vol. 807, pp. 128–153, 2020. Available: <https://doi.org/10.1016/j.tcs.2019.09.027>
- [29] L. Padovani, V. T. Vasconcelos, and H. T. Vieira, “Typing liveness in multiparty communicating systems,” in *Coordination Models and Languages*, E. Kühn and R. Pugliese, Eds. Springer, 2014, pp. 147–162. Available: https://doi.org/10.1007/978-3-662-43376-8_10

- [30] M. Carbone, O. Dardha, and F. Montesi, “Progress as compositional lock-freedom,” in *Coordination Models and Languages*, E. Kühn and R. Pugliese, Eds. Springer, 2014, pp. 49–64. Available: https://doi.org/10.1007/978-3-662-43376-8_4
- [31] S. Ghilezan, J. Pantović, I. Prokić, A. Scalas, and N. Yoshida, “Precise subtyping for asynchronous multiparty sessions,” *Proc. ACM Program. Lang.*, vol. 5, 2021. Available: <https://doi.org/10.1145/3434297>
- [32] J. Lange and E. Tuosto, “Synthesising choreographies from local session types,” in *CONCUR 2012 – Concurrency Theory*, M. Koutny and I. Ulidowski, Eds. Springer, 2012, pp. 225–239. Available: https://doi.org/10.1007/978-3-642-32940-1_17
- [33] P.-M. Deniélou and N. Yoshida, “Multiparty compatibility in communicating automata: Characterisation and synthesis of global session types,” in *ICALP ’13*, F. V. Fomin, R. Freivalds, M. Kwiatkowska, and D. Peleg, Eds. Springer, 2013, pp. 174–186. Available: https://doi.org/10.1007/978-3-642-39212-2_18
- [34] L. Padovani, “Fair subtyping for multi-party session types,” *Mathematical Structures in Computer Science*, vol. 26, no. 3, pp. 424–464, 2016. Available: <https://doi.org/10.1017/S096012951400022X>

A. Classifying Fairness Notions for our Session Calculus

In Section II-B, we have presented strong and weak fairness of transitions, and of components. We have also introduced the concepts of progress and justness. In this appendix, we discuss further fairness assumptions and relate them to each other. The notions and our classifications are based on the survey [2].

The classification from [2] defines $3 \times 6 = 18$ fairness assumptions xy with $x \in \{S, W, J\}$ and $y \in \{C, G, I, Z, A, T\}$. This is a shorthand for “ x fairness of y ”, with $x \in \{\text{strong, weak, J-}\}$ and $y \in \{\text{components, groups of components, instructions, synchronisations, actions, transitions}\}$.

Strong and weak fairness were defined in Section II-B; these notions are parametrised by the concept of a task. The parameter y above governs the choice of tasks. *Fairness of transitions*, in which each transition constitutes a task, was already defined in Section II-B.

Remember that a *component* of a network expression is one of its locations, and that *comp* is a function that associates with each transition the set of one or two components that are involved in that transition. Two transitions t and u are *concurrent*, notation $t \smile u$, iff $\text{comp}(t) \cap \text{comp}(u) = \emptyset$.

In *fairness of components*, the components constitute the tasks. Component p is *enabled* in a network state iff a transition t with $p \in \text{comp}(t)$ is enabled; a path *engages* in a component p iff it contains a transition t with $p \in \text{comp}(t)$.

In *fairness of groups of components*, the tasks are the *sets* (or *groups*) of components. A group G is *enabled* in a network state iff a transition t with $G = \text{comp}(t)$ is enabled; a path *engages* in G iff it contains a transition t with $G = \text{comp}(t)$.

Next to transitions and components, there is the concept of *instructions*. Let \mathcal{I} be the set of all occurrences of subexpressions $\lambda_k!p_k$, $\lambda_k?p_k$ or μX in a network expression \mathbb{N} . These subexpressions are called *instructions*. Each transition labelled τ stems from exactly one instruction, and each transition labelled $p \rightarrow q : \lambda$ stems from exactly two instructions. This yields the function *instr*, which associates with each transition the set of one or two instructions that gave rise to it.

In *fairness of instructions*, the instructions constitute the tasks. Instruction I is *enabled* in a network state iff a transition t with $I \in \text{instr}(t)$ is enabled; a path *engages* in an instruction I iff it contains a transition t with $I \in \text{instr}(t)$.

In *fairness of synchronisations*, the tasks are the *sets* of instructions, called *synchronisations*. A synchronisation Z is *enabled* in a network state iff a transition t with $Z = \text{instr}(t)$ is enabled; a path *engages* in a synchronisation Z iff it contains a transition t with $Z = \text{instr}(t)$.

In *fairness of actions*, the tasks are the *actions*, or transition labels. An action a is *enabled* in a network state iff a transition labelled a is enabled; a path *engages* in an action a iff it contains a transition t labelled a .

For each of these notions of a task, [2] also defines *J-fairness*. A task T is said to be enabled *during* a transition u from network state \mathbb{N} to \mathbb{N}' if T is enabled in \mathbb{N} through a transition t that is concurrent with u (i.e., $t \smile u$). Task T is

said to be *continuously* enabled if it is enabled in all network states of π and during all transitions of π . Now a path π is *J-fair* if, for each suffix π' of π , each task that is continuously enabled on π' is engaged in by π' .

Besides the 18 fairness assumptions defined above, the authors of [2] also consider progress (P) and justness (J), already defined in Section II-B, as well as *full fairness* (Fu), *extreme fairness* (Ext), *probabilistic fairness* (Pr), and *strong weak fairness of instructions* (SWI). For finite-state systems (which include the networks in our session calculus) Fu, Ext and Pr coincide with ST [2]. For this reason, there is no need to define these concepts here. Regarding SWI, say that an instruction I is *requested* in a network state \mathbb{N} if it is enabled in one of the treads in \mathbb{N} , even if it not enabled by \mathbb{N} itself, due to lack of a synchronisation partner. Now a path π is *SWI-fair* if, for each suffix π' of π , each instruction that is perpetually requested and relentlessly enabled on π' is engaged in by π' .

The following properties from [2] trivially hold (for a given network \mathbb{N}).

- (1) For each synchronisation $Z \subseteq \mathcal{I}$, and for each network state \mathbb{N} , there is at most one transition t with $\text{instr}(t) = Z$ that is enabled in \mathbb{N} .
- (2) \mathcal{I} is finite.
- (3) There is a function $cp : \mathcal{I} \rightarrow \mathcal{C}$, where \mathcal{C} is the set of components or locations in the network, such that $\text{comp}(t) = \{cp(I) \mid I \in \text{instr}(t)\}$ for all transitions t .
- (4) If an instruction I is enabled in a state \mathbb{N} , it is also requested.
- (5) If instruction I is requested in network state \mathbb{N} and u is a transition from \mathbb{N} to \mathbb{N}' such that $cp(I) \notin \text{comp}(u)$, then I is still requested in \mathbb{N}' .
- (6) If $t \smile u$ with $\text{source}(t) = \text{source}(u)$, then $\exists v \in Tr$ with $\text{source}(v) = \text{target}(u)$ and $\text{instr}(v) = \text{instr}(t)$.

Given this, the classification of fairness assumptions from [2] applies to the current setting as well, although some of these assumptions could coincide. The resulting lattice is shown in Figure 6. Here the numbers on the edges refer to the above conditions, when these are needed for the indicated comparison in strength.

When using labelling in the style of CCS, all our transitions would be labelled τ , and as a consequence, JA, WA and SA would collapse with P. But here the labelling is quite different.

A.1 SC, SA, WZ and SWI are not as strong as SI:

The following network terminates when assuming SI, for in the only infinite execution the τ -transition belonging to instruction *seller!buy* is infinitely often enabled but never taken. Termination is not guaranteed when assuming SC, SG, WZ, SWI or SA.

$$\begin{aligned} & \text{buyer} \llbracket \mu X. (\text{seller!wait}; X \oplus \text{seller!buy}) \rrbracket \\ \parallel & \text{seller} \llbracket \mu Y. (\text{buyer?wait}; Y \\ & \quad + \text{buyer?buy}; \text{shipper!order}) \rrbracket \\ \parallel & \text{shipper} \llbracket \text{seller?order} \rrbracket \end{aligned}$$

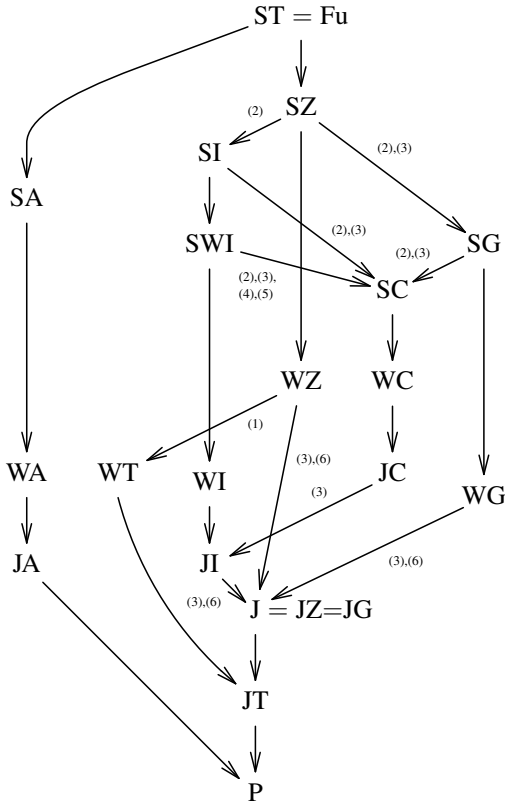


Fig. 6. A classification of progress, justness and fairness assumptions for finite-state systems [2]

A.2 WC, WG, WI, WZ and SA are not as strong as SC:

The following network terminates when assuming SC, for in any infinite execution a transition from *buyer2* is infinitely often enabled but never taken. It does not surely terminate when assuming WC, for this transition is not perpetually enabled, due to the τ -transitions of *seller*. Neither is termination guaranteed when assuming WG, WI or WZ. Furthermore, it does not surely terminate when assuming SA, because *buyer2* may be stuck before doing its initial τ -transition.

$$\begin{aligned}
& \text{seller} \llbracket \mu X. (\text{buyer1?order1}; X + \text{buyer2?order2}) \rrbracket \\
& \parallel \text{buyer1} \llbracket \mu Y. \text{seller!order1}; Y \rrbracket \\
& \parallel \text{buyer2} \llbracket \text{seller!order2} \rrbracket
\end{aligned} \quad (2)$$

A.3 Collapsing Fairness Assumptions:

Proposition 5: SG and SC coincide. WG is weaker than WC.

Proof: Let π be an infinite path in our network that is not SG-fair. One case is that an interaction between p and q is infinitely often enabled, but never taken. By taking a suffix, we may assume this interaction is enabled in the first state of π . W.l.o.g., let p be the sending party. Then process p must be in a state of the form $q!\lambda; P$, and it remains in that state for the rest of π . It follows that also component p is infinitely often enabled, but never taken. Hence π is not SC-fair.

The other case is that a single-component task (thus consisting of τ -transitions) is infinitely often enabled, but never taken. Also in this case it follows that π is not SC-fair.

The other statement is obtained in the same way. \square

Proposition 6: SZ and SI coincide. WZ is weaker than WI.

Proof: Let π be an infinite path in our network that is not SZ-fair. One case is that a synchronisation between p and q is infinitely often enabled, but never taken. By taking a suffix, we may assume this interaction is enabled in the first state of π . W.l.o.g., let p be the sending party. Then process p must be in a state of the form $q!\lambda; P$, and it remains in that state for the rest of π . It follows that also this specific instruction of p is infinitely often enabled, but never taken. Hence π is not SI-fair.

The other case is that a single-instruction task (thus consisting of τ -transitions) is infinitely often enabled, but never taken. Also in this case it follows that π is not SI-fair.

The other statement is obtained in the same way. \square

Proposition 7: SA is weaker than SC. WA is weaker than WC.

Proof: Let π be an infinite path in our network that is not SA-fair. So on π a transition label α is infinitely often enabled, but never taken. Then $\alpha \neq \tau$, because it is easy to show that each infinite path contains infinitely many τ -transitions; in fact, on any path from a network state τ -transitions make for at least half of all transitions. So α has the form $p \rightarrow q;\lambda$. In the first state of π on which α is enabled, the process p must be in a state of the form $q!\lambda; P$, and it remains in that state for the rest of π . For simplicity, we may assume that α is enabled in the first state of π ; otherwise we simply take a suffix. Hence the instruction $q!\lambda$ is perpetually requested, yet never taken. Moreover, since α is infinitely often enabled, so is component p . Yet no action from this component occurs on π . Hence π is not SC-fair.

The other statement is obtained in the same way. \square

Proposition 8: WC coincides with J (and thus also with JC and JI).

Proof: Let π be an infinite path in our network that is not WC-fair. So on π a component p is perpetually enabled, but never taken. In case p is stuck in a state where its next transition is a τ , then π is not just.

In case p is stuck in a state $q!\lambda; T$, then, for component p to be perpetually enabled, q must always be in a state $\sum_{i \in I} p_i? \lambda_i; T_i$ with $p = p_k$ and $\lambda = \lambda_k$ for some $k \in I$. Also process q must get stuck in such a state, for if q keeps moving, it will at some point reach a state $\mu X.U$, which is not of the above form. Consequently, π is not just.

The remaining case is that p is stuck in a state of the form $\sum_{i \in I} p_i? \lambda_i; T_i$. For component p to be enabled, a component p_k with $k \in I$ must be in a state $p!\lambda_k; T$. Again it follows that π is not just. \square

Proposition 9: WI coincides with J.

Proof: Let π be an infinite path in our network that is not WI-fair. So on π an instruction $\lambda!q$, $\lambda?q$ or μX from a process p is perpetually enabled, but never taken. In case of an instruction μX , π is not just.

In case of an instruction $\lambda!q$, where the τ -transition belonging to this transition is never taken, p must be stuck in a state $\bigoplus_{i \in I} p_i! \lambda_i; T_i$ with $q = p_k$ and $\lambda = \lambda_k$ for some $k \in I$; for if p performed one of the other branches, the instruction would (temporarily) cease to be enabled. Again π is not just.

In case p is stuck in a state $q! \lambda; P$, then, for that instruction $\lambda!q$ to be perpetually enabled, q must always be in a state $\sum_{i \in I} p_i? \lambda_i; T_i$ with $p = p_k$ and $\lambda = \lambda_k$ for some $k \in I$. Also process q must get stuck in such a state, for if q keeps moving, it will at some point reach a state $\mu X.U$, which is not of the above form. Consequently, π is not just.

The remaining case is that p is stuck in a state of the form $\sum_{i \in I} p_i? \lambda_i; T_i$. For component p to be enabled, a component p_k with $k \in I$ must be in a state $p! \lambda_k; T$. Again it follows that π is not just. \square

Proposition 10: WA coincides with JA.

Proof: Let π be an infinite path in our network that is not WA-fair. So on π (possibly after taking a suffix) a transition label α is perpetually enabled, but never taken. Then $\alpha \neq \tau$, as in the proof of Proposition 7. So α has the form $p \rightarrow q: \lambda$. In the first state of π , the process p must be in a state of the form $q! \lambda; P$, and it remains in that state for the rest of π . Since α is perpetually enabled, q must always be in a state $\sum_{i \in I} p_i? \lambda_i; T_i$ with $p = p_k$ and $\lambda = \lambda_k$ for some $k \in I$. Also process q must get stuck in such a state, for if q keeps moving, it will at some point reach a state $\mu X.U$, which is not of the above form. Consequently, the action α is continuously enabled on π , and π is not JA-fair. \square

Proposition 11: WT coincides with P.

Proof: Since our syntax does not allow self-loops (considering that unfolding recursion takes a τ -transition) on no infinite path a transition can be perpetually enabled. \square

Proposition 12: SWI coincides with SC.

Proof: Let π be an infinite path in our network that is not SWI-fair. So on π an instruction $\lambda!q$, $\lambda?q$ or μX from a process p is perpetually requested and infinitely often enabled, but never taken. In case of an instruction μX , π is not just, and thus certainly not SC-fair.

In case of an instruction $\lambda!q$, where the τ -transition belonging to this transition is never taken, p must be stuck in a state $\bigoplus_{i \in I} p_i! \lambda_i; T_i$ with $q = p_k$ and $\lambda = \lambda_k$ for some $k \in I$; for if p performed one of the other branches, the instruction would (temporarily) cease to be enabled. Again π is not just.

If p is stuck in a state $q! \lambda; P$, then component p is infinitely often enabled, but never taken. Hence π is not SC-fair.

In case of an instruction $\lambda?q$, p must be stuck in a state $\sum_{i \in I} p_i? \lambda_i; T_i$ with $q = p_k$ and $\lambda = \lambda_k$ for some $k \in I$; if it leaves this state, it reaches a state in which that very instruction




Fig. 7. A classification of fairness assumptions for our session calculus

$\lambda?q$ is no longer requested. Again component p is infinitely often enabled, but never taken. Hence π is not SC-fair. \square

A.4 P, WT and SA are not as strong as J: In the following network the accountant terminates when assuming J, but when merely assuming P or WT this is not guaranteed, since there is no single *accountant*-transition that is perpetually enabled. It does not surely terminate when assuming SA either, because *accountant* may be stuck before doing its initial τ -transition.

$$\begin{aligned} & \text{trader1} \llbracket \mu X. \text{trader2!order}; \text{trader2?order}; X \rrbracket \\ & \parallel \text{trader2} \llbracket \mu Y. \text{trader1?order}; \text{trader1!order}; Y \rrbracket \\ & \parallel \text{accountant} \llbracket \text{auditor!report} \rrbracket \\ & \parallel \text{auditor} \llbracket \text{accountant?report} \rrbracket \end{aligned}$$

The following network shows exactly the same, due to the initial τ -transition of *buyer*.

$$\begin{aligned} & \text{seller} \llbracket \mu X. \text{buyer?order}; X \rrbracket \\ & \parallel \text{buyer} \llbracket \mu Y. \text{seller!order}; Y \rrbracket \\ & \parallel \text{accountant} \llbracket \text{auditor!report} \rrbracket \\ & \parallel \text{auditor} \llbracket \text{accountant?report} \rrbracket \end{aligned} \quad (3)$$

A.5 P is not as strong as JA: In Network (3), the path in which the τ -action of *accountant* occurs, but the *report*-action does not, is progressing, but not JA-fair.

A.6 J and WA are not as strong as SA: In Network (2), the path in which the τ -action of *buyer2* occurs, but the *order2*-action does not, is just, as well as WA-fair, but not SA-fair.

A.7 SZ is not as strong as ST: The network in Example 3 from the introduction has 4 states and 8 transitions. ST insists that in a fair run each of these transitions occurs, whereas SZ allows a run that alternates regularly between a *buyer1/seller1*- and an *buyer2/seller2*-interaction.

It follows that our classification collapses as indicated in Figure 7.




Fig. 8. A classification of liveness properties

B. Collapsing notions of lock-freedom

Proposition 13: $\mathcal{L}(\text{SA})$ coincides with $\mathcal{L}(\text{P})$.

Proof: Using the results depicted in Figure 7 and Proposition 2, $\mathcal{L}(\text{P}) \Rightarrow \mathcal{L}(\text{SA})$.

Hence it suffices to show that if a network has a progressing path π that lacks the property of Definition 4, then it has an SA-fair path ρ that lacks this property. In case π is finite, we choose ρ as π .

In case π is infinite, we define the SA-unfairness of π as the number of different labels α such that label α is infinitely often enabled on π , but from some point onwards never taken. This must be a finite number, and if it is 0 then π is SA-fair. It now suffices to show that if the SA-unfairness of π is positive, then we can modify π into a path π' whose SA-unfairness is strictly smaller, and that still lacks the property of Definition 4.

Let α be infinitely often enabled on π , but from some point onwards never taken. As pointed out in the proof of Proposition 7, $\alpha \neq \tau$. So α has the form $p \rightarrow q : \lambda$. In the first state of π on which α is enabled, but not taken past that state, the process p must be in a state of the form $q! \lambda; P$, and it remains in that state for the rest of π . Now π can be modified into π' by skipping the last τ -transition belonging to the instruction $q! \lambda; P$ of p . This strictly decreases its SA-unfairness.

Since π fails the property of Definition 4, there must be a location p such that p does not terminate on π , and π contains only finitely many transitions that stem from component p . Now p does not terminate on π' either, and also π' contains only finitely many transitions that stem from component p . \square

Proposition 14: $\mathcal{L}(\text{ST})$ coincides with $\mathcal{L}(\text{SI})$.

Proof sketch: Proposition 2 implies $\mathcal{L}(\text{SI}) \Rightarrow \mathcal{L}(\text{ST})$.

For the other direction it suffices to show that any SI-fair path π that lacks the property of Definition 4 can be converted into an ST-fair path ρ that lacks this property. This can be achieved by swapping concurrent transitions. \blacksquare

Consequently, the 7 different fairness assumptions collapse to 4 different liveness properties, displayed in Figure 8.

C. Padovani's lock-freedom coincides with $\mathcal{L}(\text{ST})$

This appendix contains a proof of Theorem 1: a network is Padovani lock-free iff it satisfies $\mathcal{L}(\text{ST})$.

Proof: Suppose $\mathbb{N} \models \mathcal{L}(\text{ST})$. We show it is lock-free. Let \mathbb{M} be a reachable state of \mathbb{N} . Take a path from \mathbb{N} to \mathbb{M} , and extend it to an ST-fair path π . This is possible by Theorem 6.1 of [2], saying that ST-fairness is feasible. The suffix π' of π starting at \mathbb{M} satisfies the property required by Padovani's lock-freedom (Definition 5), that is, for a given location p of \mathbb{M} such that $\text{PROC}(p, \mathbb{M}) \neq \text{OK}$, π' contains a transition involving p . Interestingly, the choice of the path required by Definition 5 turns out to be independent of p .

Now suppose \mathbb{N} is lock-free, and let π be an ST-fair path. Let p be a location of \mathbb{N} such that π does not contain a state of the form $\mathbb{N}' \parallel p \llbracket \text{OK} \rrbracket$. We need to show that π contains infinitely many transitions that stems from component p .

For each state \mathbb{M} on π , let $d(\mathbb{M}) > 0$ be the length of the shortest path from \mathbb{M} that contains a transition from component p ; such a shortest path exists since \mathbb{N} is lock-free. Since \mathbb{N} is a finite-state system, there is a state \mathbb{M} that occurs infinitely often on π . In case $d(\mathbb{M}) > 1$ there must be a transition $\mathbb{M} \xrightarrow{\pi} \mathbb{M}'$ such that $d(\mathbb{M}') < d(\mathbb{M})$. Since this transition is enabled on π infinitely often, and π is ST-fair, this transition must be taken infinitely often, and hence also \mathbb{M}' occurs infinitely often in π' . So by a trivial induction there is a state \mathbb{M}'' with $d(\mathbb{M}'') = 1$ that occurs infinitely often in π . This state has an outgoing transition that stems from component p . Since this transition is enabled on π infinitely often, and π is ST-fair, it must be taken infinitely often. \square

D. Proof of completeness

This appendix contains a proof of Theorem 2: if $\mathbb{N} \models \mathcal{L}(\text{J})$, then \mathbb{N} is guardedly well-typed.

Proof: The proof is staged as a series of claims.

Claim 1: Let \mathbb{M} be unfolded.⁶ If $\mathbb{N} \xrightarrow{p \rightarrow q : \lambda} \mathbb{M}$ then $\text{PROC}(q, \mathbb{N}) = \sum_{i \in I} p_i ? \lambda_i; U_i$ with $p = p_j$, $\text{PROC}(q, \mathbb{M}) = U_j$ and $\lambda = \lambda_j$ for some $j \in I$. Moreover, if $r \notin \{p, q\}$ then $\text{PROC}(r, \mathbb{N}) = \text{PROC}(r, \mathbb{M})$.

Proof: Directly from the definition of the reactive semantics (Figure 3). \blacksquare

Claim 2: The algorithm GT (Figure 5) always terminates.

Proof: In a run on which GT does not terminate, along at least one branch an unbounded history h is created. Since there are only finitely many reachable states, some state \mathbb{M} must occur unboundedly in h . Each time this state is encountered, except for the first time, the fifth clause of GT applies. However, each time a different location p that did not already occur in $h \upharpoonright \mathbb{M}$ is added to h . Since there are only finitely many locations, this cannot go on forever. \blacksquare

An expression $\text{GT}(h, \mathbb{N})$ may have free occurrences of variables $X_{\mathbb{M}}$. It is easy to check that if $X_{\mathbb{M}}$ has a free occurrence in $\text{GT}(h, \mathbb{N})$, then \mathbb{M} occurs in h . We define a

⁶See Page 9 for a definition.

closed version $GT^*(h, \mathbb{N})$ of $GT(h, \mathbb{N})$, obtained from $GT(h, \mathbb{N})$ by unfolding recursion. The definition proceeds by induction on the length of h . Here, $\text{fv}(\mathcal{G})$ denotes the set of free recursion variables in a global type expression \mathcal{G} .

$$GT^*(h, \mathbb{N}) := GT(h, \mathbb{N}) \left\{ GT^*(h \upharpoonright \mathbb{M}, \mathbb{M}) /_{X_{\mathbb{M}}} \left[X_{\mathbb{M}} \in \text{fv}(GT(h, \mathbb{N})) \right] \right\}$$

Note that $GT(\varepsilon, \mathbb{N}) = GT^*(\varepsilon, \mathbb{N})$. By induction, $GT^*(h, \mathbb{N})$ is a closed session type expression.

$GT(h, \mathbb{N})$, and hence also $GT^*(h, \mathbb{N})$, always yields a valid global session type expression, except that it may contain the constant DEADLOCK.

Claim 3: If $GT(h, \mathbb{N})$ contains the constant DEADLOCK, then \mathbb{N} is not deadlock-free.

Proof: If $GT(h, \mathbb{N})$ contains the constant DEADLOCK, then for some unfolded network \mathbb{M} reachable from \mathbb{N} and for some extension h' of h we have $GT(h', \mathbb{M}) = \text{DEADLOCK}$. It suffices to show that \mathbb{M} has a deadlock. Since no location is ready in \mathbb{M} , for each location p of \mathbb{M} with $\text{PROC}(p, \mathbb{M}) = \bigoplus_{i \in I} q_i! \lambda_i; \mathbf{T}_i$ – let us call such a location *active* in \mathbb{M} – there exists an $i \in I$ such that there is no transition $\mathbb{M} \xrightarrow{p \rightarrow q_i: \lambda_i} \mathbb{M}_i$. Let i_p be this i . Now \mathbb{M} admits a sequence of τ -transitions to a state \mathbb{M}' in which $\text{PROC}(p, \mathbb{M}') = \ulcorner q_{i_p}! \lambda_{i_p}; \mathbf{T}_{i_p}$ for each p active in \mathbb{M} , and no further τ -transitions are possible from \mathbb{M}' . The only transitions that \mathbb{M}' could possibly do must have a label $p \rightarrow q_{i_p}: \lambda_{i_p}$ for some p active in \mathbb{M} , yet none of these transitions are actually possible. Hence \mathbb{M}' is a deadlock. ■

Call a history h *reachable* from \mathbb{N} iff all networks \mathbb{M} that occur in h are reachable from \mathbb{N} .

Claim 4: If \mathbb{N} is deadlock-free, and h and \mathbb{M} are reachable from \mathbb{N} , then $GT^*(h, \mathbb{M})$ does not contain DEADLOCK.

Proof: If $GT^*(h, \mathbb{M})$ contains the constant DEADLOCK, then either $GT(h, \mathbb{M})$ contains DEADLOCK, or $GT^*(h', \mathbb{M}')$ does, for a proper prefix h' of h and a network \mathbb{M}' that occurs in h . The previous claim and a simple induction on the length of h finish the proof. ■

Claim 5: Let $\mathbb{N} \models \mathcal{L}(\mathbf{J})$, let r be a location of \mathbb{N} , and h and \mathbb{M} be reachable from \mathbb{N} , with \mathbb{M} unfolded.

If $r \notin \text{parties}(GT^*(h, \mathbb{M}))$, then $\text{PROC}(r, \mathbb{M}) = \text{OK}$.

Proof: For each pair (ℓ, \mathbb{L}) of a history ℓ and a network state \mathbb{L} , both reachable from \mathbb{N} , such that $\text{PROC}(r, \mathbb{L}) \neq \text{OK}$, we select a unique successor pair (ℓ', \mathbb{L}') as follows, inspired by the definition of $GT(\ell, \mathbb{L})$. In case \mathbb{L} is not unfolded, we pick a network \mathbb{L}' with $\mathbb{L} \xrightarrow{\tau} \mathbb{L}'$ and take $\ell' := \ell$. Otherwise, in case \mathbb{L} does not occur in ℓ or $\ell \upharpoonright \mathbb{L}$ is incomplete for \mathbb{L} , let $p := \text{CH}(\ell, \mathbb{L})$ and $\text{PROC}(p, \mathbb{L}) = \bigoplus_{i \in I} q_i! \lambda_i; \mathbf{T}_i$. Now pick a $k \in I$ and take $\ell' := \ell_k$ and $\mathbb{L}' := \mathbb{L}_k^p$ (as in the definition of $GT(\ell, \mathbb{L})$). Finally, if \mathbb{L} is unfolded, \mathbb{L} occurs in ℓ , and $\ell \upharpoonright \mathbb{L}$ is complete for \mathbb{L} , we take the unique successor pair of $(\ell \upharpoonright \mathbb{L}, \mathbb{L})$.

If (ℓ', \mathbb{L}') is the successor of (ℓ, \mathbb{L}) then surely there is a transition $\mathbb{L} \xrightarrow{\tau} \mathbb{L}'$ or $\mathbb{L} \xrightarrow{p \rightarrow q_k: \lambda_k} \mathbb{L}'$ with $p = \text{CH}(\ell, \mathbb{L})$ or $p = \text{CH}(\ell \upharpoonright \mathbb{L}, \mathbb{L})$ and k as chosen above.⁷ Combining those

⁷Each transition $\xrightarrow{p \rightarrow q: \lambda}$ can be split into two transitions $\xrightarrow{\tau} \xrightarrow{p \rightarrow q: \lambda}$.

transitions yields for each pair (ℓ, \mathbb{L}) a unique path $\pi(\ell, \mathbb{L})$ starting from \mathbb{L} , which is either infinite or ends in a state \mathbb{L}' with $\text{PROC}(r, \mathbb{L}') = \text{OK}$. Here we use Claim 4. Moreover, in case $\pi(\ell, \mathbb{L})$ is infinite, by (the proof of) Claim 2 it must have a suffix $\pi(\ell', \mathbb{L}')$ such that \mathbb{L}' is an unfolded network expression, \mathbb{L}' occurs in ℓ' , and $\ell' \upharpoonright \mathbb{L}'$ is complete for \mathbb{L}' . Hence that suffix is a simple loop. By construction, the pair (ℓ', \mathbb{L}') must be unique; call $\pi(\ell', \mathbb{L}')$ the *loop suffix* of $\pi(\ell, \mathbb{L})$.

First assume that $\pi(h, \mathbb{M})$ does contain a transition that involves component r . Considering that \mathbb{M} is unfolded, we have $\text{PROC}(r, \mathbb{M}) \neq \mu X.T$. Hence, this transition must have the label $p \rightarrow q: \lambda$ with $p = r$ or $q = r$. A simple induction shows that $r \in \text{parties}(GT^*(h, \mathbb{M}))$.

Henceforth, we assume that $\pi(h, \mathbb{M})$ contains no transition involving component r . First assume that $\pi(h, \mathbb{M})$ ends in a state \mathbb{M}' with $\text{PROC}(r, \mathbb{M}') = \text{OK}$. Since $\pi(h, \mathbb{M})$ contains no transition involving r , $\text{PROC}(r, \mathbb{M}) = \text{OK}$. Finally, assume that $\pi(h, \mathbb{M})$ is infinite. It suffices to derive a contradiction.

Let $\pi(\ell, \mathbb{L})$ be the loop suffix of $\pi(h, \mathbb{M})$. Now $\pi(\ell, \mathbb{L})$ is infinite and contains no transition that involves component r . Moreover, \mathbb{L} is reachable from \mathbb{N} .

Suppose that there is a location p with $\text{PROC}(p, \mathbb{L}) = \bigoplus_{i \in I} q_i! \lambda_i; \mathbf{T}_i$, no transition in $\pi(\ell, \mathbb{L})$ involves component p from one of the components q_i for $i \in I$, and $\mathbb{L} \xrightarrow{p \rightarrow q_i: \lambda_i} \mathbb{L}'$ for all $i \in I$. In that case p is ready in \mathbb{L} , and $\ell \upharpoonright \mathbb{L}$ is incomplete for \mathbb{L} . This contradicts the definition of the loop suffix.

It follows that for each location p with $\text{PROC}(p, \mathbb{L}) = \bigoplus_{i \in I} q_i! \lambda_i; \mathbf{T}_i$, and such that no transition in $\pi(\ell, \mathbb{L})$ involves component p , there exists an $k \in I$ such that either infinitely many transitions in $\pi(\ell, \mathbb{L})$ involve component q_k , or $\mathbb{L} \xrightarrow{p \rightarrow q_k: \lambda_k} \mathbb{L}'$. Let π' be the infinite path obtained from $\pi(\ell, \mathbb{L})$ by transforming all states \mathbb{L}' in this path in the same way, namely by replacing, for all locations p as above, $\text{PROC}(p, \mathbb{L}) = \bigoplus_{i \in I} q_i! \lambda_i; \mathbf{T}_i$ by the appropriate $\ulcorner q_k! \lambda_k; \mathbf{T}_k$. By construction this path is just.

The path π' is the suffix of a path π'' that starts in \mathbb{N} . This path contains only finitely many transitions that involve component r , and no state $\mathbb{N}' \parallel r \llbracket \text{OK} \rrbracket$. Consequently, \mathbb{N} does not satisfy $\mathcal{L}(\mathbf{J})$. ■

Now assume that \mathbb{N} satisfies $\mathcal{L}(\mathbf{J})$. Then \mathbb{N} is deadlock-free, and hence $GT(\varepsilon, \mathbb{N})$ yields a valid global session type, by Claim 3. To prove that \mathbb{N} is well-typed w.r.t. $GT(\varepsilon, \mathbb{N})$, it suffices to show that $\mathbb{N} \vdash GT(\varepsilon, \mathbb{N})$. In fact, we prove a stronger claim, namely that for all histories h and networks \mathbb{M} that are both reachable from \mathbb{N} we have $\mathbb{M} \vdash GT^*(h, \mathbb{M})$.

By construction, $\text{parties}(GT(h, \mathbb{M}))$ contains only locations of \mathbb{M} , and hence of \mathbb{N} . Therefore, the same holds for $\text{parties}(GT^*(h, \mathbb{M}))$. Thus, it remains to establish that $\text{PROC}(r, \mathbb{M}) \vdash GT^*(h, \mathbb{M}) \upharpoonright_r$ for all h and \mathbb{M} reachable from \mathbb{N} , and all locations r of \mathbb{N} . We do this by coinduction [19]. We make a case distinction on the shape of $GT(h, \mathbb{M})$, and apply induction on h , and a nested induction on the number of recursion-unfolding $\xrightarrow{\tau}$ -transitions possible from \mathbb{M} . Pick h , \mathbb{M} and r in the following.

- Suppose that $\mathbb{M} \xrightarrow{\tau} \mathbb{M}'$ for a network \mathbb{M}' , that is, $\text{PROC}(p, \mathbb{M}) = \mu X.T$ and $\text{PROC}(p, \mathbb{M}') = T\{\mu X.T/X\}$ for some location p of \mathbb{N} .

In case $r=p$, using the first rule for \vdash , we derive $\text{PROC}(r, \mathbb{M}) = \mu X.T \vdash \text{GT}^*(h, \mathbb{M}) \upharpoonright_r$ from $P\{\mu X.T/X\} = \text{PROC}(r, \mathbb{M}') \vdash \text{GT}^*(h, \mathbb{M}') \upharpoonright_r = \text{GT}^*(h, \mathbb{M}) \upharpoonright_r$, and the latter is a coinduction hypothesis.

In case $r \neq p$, then $\text{PROC}(r, \mathbb{M}) = \text{PROC}(r, \mathbb{M}')$ and $\text{GT}^*(h, \mathbb{M}) \upharpoonright_r = \text{GT}^*(h, \mathbb{M}') \upharpoonright_r$. Since \mathbb{M}' admits fewer recursion-unfolding $\xrightarrow{\tau}$ -transitions than \mathbb{M} , by induction, $\text{PROC}(r, \mathbb{M}') \vdash \text{GT}^*(h, \mathbb{M}') \upharpoonright_r$.

In the remainder, we assume that \mathbb{M} is already unfolded.

- Let $\text{GT}(h, \mathbb{M}) = \text{OK}$. Then $\text{GT}^*(h, \mathbb{M}) = \text{PROC}(r, \mathbb{M}) = \text{OK}$. Now the third rule for \vdash yields $\text{PROC}(r, \mathbb{M}) \vdash \text{GT}^*(h, \mathbb{M}) \upharpoonright_r$.
- Let $\text{GT}(h, \mathbb{M}) = X_{\mathbb{M}}$. Then \mathbb{M} occurs in h and we have $\text{GT}^*(h, \mathbb{M}) = \text{GT}^*(h \upharpoonright \mathbb{M}, \mathbb{M})$. By induction, since $h \upharpoonright \mathbb{M}$ is strictly shorter than h , $\text{PROC}(r, \mathbb{M}) \vdash \text{GT}^*(h, \mathbb{M}) \upharpoonright_r$.
- Let $\text{GT}(h, \mathbb{M}) = \boxplus_{i \in I} p \rightarrow q_i : \lambda_i ; \text{GT}(h_i, \mathbb{M}_i^p)$, so $\text{GT}^*(h, \mathbb{M}) = \boxplus_{i \in I} p \rightarrow q_i : \lambda_i ; \text{GT}^*(h_i, \mathbb{M}_i^p)$. If $p=r$, then $\text{GT}^*(h, \mathbb{M}) \upharpoonright_r = \bigoplus_{i \in I} q_i ! \lambda_i ; (\text{GT}^*(h_i, \mathbb{M}_i^p) \upharpoonright_r)$. By the coinduction hypothesis, for each $i \in I$ we may assume $\text{PROC}(p, \mathbb{M}_i^p) \vdash \text{GT}^*(h_i, \mathbb{M}_i^p) \upharpoonright_p$. Moreover, $\text{PROC}(p, \mathbb{M}) = \bigoplus_{i \in I} q_i ! \lambda_i ; T_i$ with $T_i = \text{PROC}(p, \mathbb{M}_i^p)$. Now apply the fifth proof rule for \vdash . If $p \neq r$, $\text{GT}^*(h, \mathbb{M}) \upharpoonright_r = \prod_{i \in I} (p \rightarrow q_i : \lambda_i ; \text{GT}(h_i, \mathbb{M}_i^p)) \upharpoonright_r$. Pick $k \in I$. Applying the last rule for \vdash , we need to show that $\text{PROC}(r, \mathbb{M}) \vdash (p \rightarrow q_k : \lambda_k ; \text{GT}(h_k, \mathbb{M}_k^p)) \upharpoonright_r$. If $r \neq q_k$, $(p \rightarrow q_k : \lambda_k ; \text{GT}(h_k, \mathbb{M}_k^p)) \upharpoonright_r = \text{GT}(h_k, \mathbb{M}_k^p) \upharpoonright_r$, and $\text{PROC}(r, \mathbb{M}) = \text{PROC}(r, \mathbb{M}_k^p)$ by Claim 1. Moreover, $\text{PROC}(r, \mathbb{M}_k^p) \vdash \text{GT}(h_k, \mathbb{M}_k^p) \upharpoonright_r$ can be used as coinduction hypothesis. If $r = q_k$ then

$$(p \rightarrow q_k : \lambda_k ; \text{GT}(h_k, \mathbb{M}_k^p)) \upharpoonright_r = p ? \lambda_k ; \text{GT}(h_k, \mathbb{M}_k^p) \upharpoonright_r.$$

Moreover, by Claim 1, $\text{PROC}(r, \mathbb{M}) = \sum_{h \in H} p h ? \lambda_h ; U_h$ with $p = p_j$, $\text{PROC}(r, \mathbb{M}_k^p) = U_h$ and $\lambda_k = \lambda_j$ for some $j \in H$. Using $\text{PROC}(r, \mathbb{M}_k^p) \vdash \text{GT}(h_k, \mathbb{M}_k^p)$ as coinduction hypothesis, $\text{PROC}(r, \mathbb{M}) \vdash p ? \lambda_k ; \text{GT}(h_k, \mathbb{M}_k^p) \upharpoonright_r$ follows by application of the fourth rule for \vdash .

- Let $\text{GT}(h, \mathbb{M}) = \mu X_{\mathbb{M}}. \boxplus_{i \in I} p \rightarrow q_i : \lambda_i ; \text{GT}(h_i, \mathbb{M}_i^p)$. Let $\text{GT}^*(h, \mathbb{N}, X_{\mathbb{M}})$ be defined as $\text{GT}^*(h, \mathbb{N})$, except that the free variable $X_{\mathbb{M}}$ does not get unfolded. Then $\text{GT}^*(h, \mathbb{N}) = \text{GT}^*(h, \mathbb{N}, X_{\mathbb{M}})\{\text{GT}^*(h \upharpoonright \mathbb{M}, \mathbb{M})/X_{\mathbb{M}}\}$. Hence $\text{GT}^*(h, \mathbb{M}) = \mu X. \boxplus_{i \in I} p \rightarrow q_i : \lambda_i ; \text{GT}^*(h_i, \mathbb{M}_i^p, X_{\mathbb{M}})$. First let $r \notin \text{parties}(\text{GT}^*(h, \mathbb{M}))$. Then $\text{GT}^*(h, \mathbb{M}) \upharpoonright_r = \text{OK}$. By Claim 5, $\text{PROC}(r, \mathbb{M}) = \text{OK}$. Consequently, $\text{PROC}(r, \mathbb{M}) \vdash \text{GT}^*(h, \mathbb{M}) \upharpoonright_r$ via the third rule for \vdash . Now if $r \in \text{parties}(\text{GT}^*(h, \mathbb{M}))$ then $\text{GT}^*(h, \mathbb{M}) \upharpoonright_r = \mu X. ((\boxplus_{i \in I} p \rightarrow q_i : \lambda_i ; \text{GT}^*(h_i, \mathbb{M}_i^p, X_{\mathbb{M}})) \upharpoonright_r)$. Since \mathbb{M} does not occur in h we have $h_i \upharpoonright \mathbb{M} = h$. Hence

$$\begin{aligned} \text{GT}^*(h_i, \mathbb{M}_i^p) \upharpoonright_r &= (\text{GT}^*(h_i, \mathbb{M}_i^p, X_{\mathbb{M}})\{\text{GT}^*(h \upharpoonright \mathbb{M}, \mathbb{M})/X_{\mathbb{M}}\}) \upharpoonright_r \\ &= (\text{GT}^*(h_i, \mathbb{M}_i^p, X_{\mathbb{M}})\{\text{GT}^*(h, \mathbb{M})/X_{\mathbb{M}}\}) \upharpoonright_r \\ &= \text{GT}^*(h_i, \mathbb{M}_i^p, X_{\mathbb{M}}) \upharpoonright_r \{\text{GT}^*(h, \mathbb{M}) \upharpoonright_r / X_{\mathbb{M}}\}. \end{aligned}$$

In order to obtain $\text{PROC}(r, \mathbb{M}) \vdash \text{GT}^*(h, \mathbb{M}) \upharpoonright_r$, by the second rule for \vdash it suffices to establish $\text{PROC}(r, \mathbb{M}) \vdash$

$$\begin{aligned} &(\boxplus_{i \in I} p \rightarrow q_i : \lambda_i ; \text{GT}^*(h_i, \mathbb{M}_i^p, X_{\mathbb{M}})) \upharpoonright_r \{\text{GT}^*(h, \mathbb{M}) \upharpoonright_r / X_{\mathbb{M}}\} \\ &= (\boxplus_{i \in I} p \rightarrow q_i : \lambda_i ; \text{GT}^*(h_i, \mathbb{M}_i^p)) \upharpoonright_r. \end{aligned}$$

This proceeds exactly as in the previous case.

This shows that \mathbb{N} is well-typed w.r.t. $\text{GT}(\varepsilon, \mathbb{N})$. It remains to show that all projections $\text{GT}(\varepsilon, \mathbb{N}) \upharpoonright_p$ are guarded.

Claim 6: Let $\mathbb{N} \models \mathcal{L}(\mathbb{J})$ and let r be a location of \mathbb{N} . Then $\text{GT}(\varepsilon, \mathbb{N}) \upharpoonright_r$ is guarded, i.e., it does not have a subexpression of the form $\mu X.U$ such that X occurs in U outside the scope of any subexpression $p ! \lambda ; T$ or $p ? \lambda ; T$.

Proof: Suppose, towards a contradiction, that $\text{GT}(\varepsilon, \mathbb{N}) \upharpoonright_r$ does have a subexpression of the form $\mu X.U$ such that X occurs in U outside the scope of any subexpression $p ! \lambda ; T$ or $p ? \lambda ; T$. By the definition of projection, this subexpression must have the form $\mathcal{G} \upharpoonright_r$, with $\mathcal{G} = \mu X.\mathcal{G}'$ a subexpression of $\text{GT}(\varepsilon, \mathbb{N})$. Given the algorithm of Figure 5, \mathcal{G} must have the form $\text{GT}(h, \mathbb{M})$ for a history h and network \mathbb{M} reachable from \mathbb{N} . Moreover, $X = X_{\mathbb{M}}$.

There must be a path in the parse tree of U towards the unguarded occurrence of X . Since the occurrence is unguarded, this path passes only through operators μY and $\prod_{i \in I}$. Backtracking this path through the projection from $\text{GT}(h, \mathbb{M})$ yields a path ρ in the parse tree of $\text{GT}(h, \mathbb{M})$ to a subexpression $\text{GT}(\ell, \mathbb{M}) = X_{\mathbb{M}}$, with $\ell \upharpoonright \mathbb{M}$ an extension of h . This path ρ passes merely through operators $\boxplus_{i \in I} p \rightarrow q_i : \lambda_i$ with r not being among p and the q_i .

The syntactic path ρ induces a path π' in the transition system from \mathbb{M} to \mathbb{M} .

Suppose that there is a location p with $\text{PROC}(p, \mathbb{M}) = \bigoplus_{i \in I} q_i ! \lambda_i ; T_i$, no transition in π' involves component p or from one of the components q_i for $i \in I$, and $\mathbb{M} \xrightarrow{p \rightarrow q_i : \lambda_i} \rightarrow$ for all $i \in I$. In that case p is ready in \mathbb{M} , and $\ell \upharpoonright \mathbb{M}$ is incomplete for \mathbb{M} . This contradicts the definition of $\text{GT}(\ell, \mathbb{M})$.

It follows that for each location p with $\text{PROC}(p, \mathbb{M}) = \bigoplus_{i \in I} q_i ! \lambda_i ; T_i$, and such that no transition in π' involves component p , there exists a $k \in I$ such that either some transitions in π' stem from component q_k , or $\mathbb{M} \xrightarrow{p \rightarrow q_k : \lambda_k} \rightarrow$. Let π'' be the infinite path obtained from π' by transforming all states \mathbb{L} in this path in the same way, namely by replacing, for all locations p as above, $\text{PROC}(p, \mathbb{L}) = \bigoplus_{i \in I} q_i ! \lambda_i ; T_i$ by the appropriate $\lceil q_k ! \lambda_k ; T_k$.

Let π be the path from \mathbb{N} to \mathbb{M} , followed by infinitely many repetitions of the loop π'' . By construction this path is just. Past \mathbb{M} , π contains no transitions involving location r . Thus, invoking the assumption that $\mathbb{N} \models \mathcal{L}(\mathbb{J})$, it follows that r successfully terminates on π , that is, $\text{PROC}(r, \mathbb{M}) = \text{OK}$. The algorithm of Figure 5 implies that $r \notin \text{parties}(\text{GT}(h, \mathbb{M}))$. Now $\text{GT}(h, \mathbb{M}) = \mathcal{G} = \mu X.\mathcal{G}'$ is not closed, for it were, that would imply that $\mathcal{G} \upharpoonright_r = \text{OK}$, contradicting the assumption that $\mathcal{G} \upharpoonright_r = \mu X.U$. So, $\text{GT}(h, \mathbb{M})$ occurs within a subexpression $\text{GT}(h', \mathbb{L}) = \mu Y.\mathcal{H}$ of $\text{GT}(\varepsilon, \mathbb{N})$, with h' a strict prefix of h , and such that Y occurs freely in $\text{GT}(h, \mathbb{M})$. Here Y must have the form $\text{GT}(h'', \mathbb{L})$, with h'' an extension of h . Thus \mathbb{L} is

reachable from \mathbb{M} and hence $\text{PROC}(r, \mathbb{L}) = \text{OK}$. Again, it follows that $r \notin \text{parties}(\text{GT}(h', \mathbb{L}))$ and also $\text{GT}(h', \mathbb{L})$ is not closed. Going on this way, we eventually find a subexpression $\mu Z.\mathcal{H}'$ of $\text{GT}(\varepsilon, \mathbb{N})$ that is not closed, but also not inside another expression $\mu W.\mathcal{H}''$. This contradicts with $\text{GT}(\varepsilon, \mathbb{N})$ being closed. \blacksquare \square

E. Proof of Soundness

This appendix contains the proof of soundness for guardedly well-typed and race-free networks (Theorem 3).

Definition 10: To type network states we extend our type system with the following rule.

$$\frac{k \in I \quad \mathbb{T}_k \vdash \mathbb{U}_k}{\ulcorner q_k! \lambda_k; \mathbb{T}_k \vdash \bigoplus_{i \in I} q_i! \lambda_i; \mathbb{U}_i}$$

Session fidelity for recursion and internal choice can be proven independently. These lemmas show that τ -transitions preserve the type of a network.

Lemma 1: If $\mathbb{N} \xrightarrow{\tau} \mathbb{M}$ with $\text{PROC}(p, \mathbb{N}) = \mu X.T$ and $\text{PROC}(p, \mathbb{M}) = \mathbb{T}\{\mu X.T/X\}$, and $\mathbb{N} \vdash^g \mathcal{G}$ then $\mathbb{M} \vdash^g \mathcal{G}$.

Proof: If $\mathbb{N} \vdash^g \mathcal{G}$ then $\mu X.T \vdash \mathcal{G}|_p$ and $\mathcal{G}|_p$ is guarded. By the type rules, this can hold only if $\mathbb{T}\{\mu X.T/X\} \vdash \mathcal{G}|_p$. Therefore, using that $\text{PROC}(q, \mathbb{N}) = \text{PROC}(q, \mathbb{M})$ for all locations $q \neq p$, $\mathbb{M} \vdash^g \mathcal{G}$. \square

Let $\text{loc}(\mathbb{N})$ denote the set of locations of a network state \mathbb{N} .

Lemma 2: If $\mathbb{N} \xrightarrow{\tau} \mathbb{M}$ with $\text{PROC}(p, \mathbb{N}) = \bigoplus_{i \in I} q_i! \lambda_i; \mathbb{T}$, and $\text{PROC}(p, \mathbb{M}) = \ulcorner q_i! \lambda_i; \mathbb{T}_i$ for some $i \in I$ and $\mathbb{N} \vdash^g \mathcal{G}$ then $\mathbb{M} \vdash^g \mathcal{G}$.

Proof: Assume $\mathbb{N} \vdash^g \mathcal{G}$. So, \mathcal{G} is closed, $\text{parties}(\mathcal{G}) \subseteq \text{loc}(\mathcal{G})$, and, for all $p \in \text{loc}(\mathbb{N})$, $\text{PROC}(p, \mathbb{N}) \vdash \mathcal{G}|_p$ and $\mathcal{G}|_p$ is guarded. The rules of Figure 4 imply that when $\bigoplus_{i \in I} q_i! \lambda_i; \mathbb{T} \vdash \mathcal{H}$, and $i \in I$, certainly also $\ulcorner q_i! \lambda_i; \mathbb{T}_i \vdash \mathcal{H}$. As $\text{PROC}(q, \mathbb{N}) = \text{PROC}(q, \mathbb{M})$ for all locations $q \neq p$, this implies $\mathbb{M} \vdash^g \mathcal{G}$. \square

For transitions of the form $\mathbb{N} \xrightarrow{p \rightarrow q: \lambda} \mathbb{M}$ we target a session fidelity result, which is stronger than subject reduction, since it constructs a type for \mathbb{M} from the type of \mathbb{N} that reflects the network transition. For this, we need a few auxiliary concepts.

Definition 11: The *maximum depth* $\|\mathcal{G}\|_p$ in the abstract syntax tree of \mathcal{G} of a communication involving location p is defined as follows:

$$\begin{aligned} \|\text{OK}\|_p &= 0 \\ \|\mathbb{X}\|_p &= \infty \\ \|\mu X.\mathcal{G}\|_p &= \begin{cases} 0 & \text{if } p \notin \text{parties}(\mathcal{G}) \\ & \text{and } \mu X.\mathcal{G} \text{ is closed} \\ 1 + \|\mathcal{G}\|_p & \text{otherwise} \end{cases} \\ \|\boxplus_{i \in I} r \rightarrow q_i: \lambda_i; \mathcal{G}_i\|_p &= \max \left\{ \|r \rightarrow q_i: \lambda_i; \mathcal{G}_i\|_p : i \in I \right\} \\ \|r \rightarrow q: \lambda; \mathcal{G}\|_p &= \begin{cases} 1 & \text{if } p = r \vee p = q \\ 1 + \|\mathcal{G}\|_p & \text{if } p \neq r \wedge p \neq q \end{cases} \end{aligned}$$

Call a projection type \mathbb{T} *fully guarded*, if it is guarded, and each occurrence of a variable X within \mathbb{T} occurs within a subexpression $p! \lambda; \mathbb{U}$ or $p? \lambda; \mathbb{U}$ of \mathbb{T} .

Lemma 3: If $\mathcal{G}|_p$ is fully guarded then $\|\mathcal{G}\|_p$ is finite.

Proof: A straightforward structural induction on \mathcal{G} , using that

- $(\mu X.\mathcal{G})|_p$ is fully guarded iff $\mathcal{G}|_p$ is fully guarded;
- if $p \neq r, q$ then $(r \rightarrow q: \lambda; \mathcal{G})|_p = \mathcal{G}|_p$. \square

Corollary 3: If $\mathbb{N} \vdash^g \mathcal{G}$ and $p \in \text{loc}(\mathbb{N})$ then $\|\mathcal{G}\|_p$ is finite.

Proof: Let $\mathbb{N} \vdash^g \mathcal{G}$ and $p \in \text{loc}(\mathbb{N})$. By Definition 7, $\mathcal{G}|_p$ is guarded. By Definition 6, \mathcal{G} , and hence also $\mathcal{G}|_p$, is closed. Since a closed projection type is fully guarded iff it is guarded, the result follows from Lemma 3. \square

Lemma 4: If \mathcal{H} is closed then $(\mathcal{G}\{\mathcal{H}/X\})|_p = \mathcal{G}|_p\{\mathcal{H}|_p/X\}$.

Proof: A trivial structural induction on \mathcal{G} . \square

Lemma 5: $\mathbb{N} \vdash^g \mu X.\mathcal{G}$ iff $\mathbb{N} \vdash^g \mathcal{G}\{\mu X.\mathcal{G}/X\}$.

Proof: $\mu X.\mathcal{G}$ is closed iff $\mathcal{G}\{\mu X.\mathcal{G}/X\}$ is closed. Moreover, $\text{parties}(\mu X.\mathcal{G}) = \text{parties}(\mathcal{G}) = \text{parties}(\mathcal{G}\{\mu X.\mathcal{G}/X\})$. Pick $p \in \text{loc}(\mathbb{N})$. It remains to show that $(\mu X.\mathcal{G})|_p$ is guarded iff $\mathcal{G}\{\mu X.\mathcal{G}/X\}|_p$ is guarded, and $\text{PROC}(p, \mathbb{N}) \vdash (\mu X.\mathcal{G})|_p$ iff $\text{PROC}(p, \mathbb{N}) \vdash \mathcal{G}\{\mu X.\mathcal{G}/X\}|_p$.

If $p \notin \text{parties}(\mathcal{G})$ then $(\mu X.\mathcal{G})|_p = \text{OK}$ and $(\mathcal{G}\{\mu X.\mathcal{G}/X\})|_p$ must be OK in the scope of some merge operators only. Both are guarded. Moreover, $\text{PROC}(p, \mathbb{N}) \vdash \text{OK}$ iff $\text{PROC}(p, \mathbb{N}) \vdash (\mathcal{G}\{\mu X.\mathcal{G}/X\})|_p$.

If $p \in \text{parties}(\mathcal{G})$ then we have $(\mu X.\mathcal{G})|_p = \mu X.(\mathcal{G}|_p)$ and $(\mathcal{G}\{\mu X.\mathcal{G}/X\})|_p = \mathcal{G}|_p\{\mu X.(\mathcal{G}|_p)/X\}$ by Lemma 4. Now $\mathcal{G}|_p\{\mu X.(\mathcal{G}|_p)/X\}$ is guarded iff $\mu X.(\mathcal{G}|_p)$ is guarded. Moreover, by the second rule for \vdash , $\text{PROC}(p, \mathbb{N}) \vdash \mu X.(\mathcal{G}|_p)$ iff $\text{PROC}(p, \mathbb{N}) \vdash \mathcal{G}|_p\{\mu X.(\mathcal{G}|_p)/X\}$. \square

Lemma 6: If \mathcal{G} is closed then $\|\mu X.\mathcal{G}\|_p > \|\mathcal{G}\{\mu X.\mathcal{G}/X\}\|_p$ for all locations $p \in \text{parties}(\mathcal{G})$.

Proof: For any \mathcal{G} and closed \mathcal{H} we have $\|\mathcal{G}\|_p \geq \|\mathcal{G}\{\mathcal{H}/X\}\|_p$, by a trivial induction on the structure of \mathcal{G} . Hence $\|\mu X.\mathcal{G}\|_p = 1 + \|\mathcal{G}\|_p \geq 1 + \|\mathcal{G}\{\mu X.\mathcal{G}/X\}\|_p > \|\mathcal{G}\{\mu X.\mathcal{G}/X\}\|_p$. \square

Lemma 7: If $\text{PROC}(p, \mathbb{N}) \vdash \mathcal{G}|_p$ with $p \in \text{loc}(\mathbb{N}) \setminus \text{parties}(\mathcal{G})$ and \mathcal{G} is closed, then $\text{PROC}(p, \mathbb{N}) \vdash \text{OK}$.

Proof: A trivial structural induction on \mathcal{G} . \square

Let \longrightarrow be the transition relation between global session types defined in Figure 9. In combination with Lemmas 1 and 2, the following session fidelity result shows how race-free networks evolve according to the global type.

Lemma 8: For race-free network states \mathbb{N} , if $\mathbb{N} \xrightarrow{p \rightarrow q: \lambda} \mathbb{M}$ and $\mathbb{N} \vdash^g \mathcal{G}$ then there exists \mathcal{G}' such that $\mathcal{G} \xrightarrow{p \rightarrow q: \lambda} \mathcal{G}'$ and $\mathbb{M} \vdash^g \mathcal{G}'$.

Proof: By Corollary 3, $\|\mathcal{G}\|_p$ is finite. The proof proceeds by induction on $\|\mathcal{G}\|_p$. Note that $\text{PROC}(p, \mathbb{N})$ has the form $\ulcorner q! \lambda; \mathbb{T}$. Since $\ulcorner q! \lambda; \mathbb{T} \not\vdash \text{OK}$, we can rule out that $\mathcal{G} = \text{OK}$.

Let $\mathcal{G} = \mu X.\mathcal{H}$. Since $\mathbb{N} \vdash^g \mathcal{G}$, one has $\text{PROC}(p, \mathbb{N}) \vdash \mathcal{G}|_p$ and \mathcal{G} is closed. By Lemma 7, since $\ulcorner q! \lambda; \mathbb{T} \not\vdash \text{OK}$, we have $p \in \text{parties}(\mathcal{G})$. By Lemma 5, $\mathbb{N} \vdash^g \mathcal{H}\{\mu X.\mathcal{H}/X\}$. By Lemma 6,

$$\frac{\frac{k \in I}{\prod_{i \in I} p \rightarrow q_i : \lambda_i ; \mathcal{G}_i \xrightarrow{p \rightarrow q_k : \lambda_k} \mathcal{G}_k}}{\prod_{i \in J} r \rightarrow s_i : \lambda_i ; \mathcal{G}_i \xrightarrow{p \rightarrow q : \lambda} \prod_{i \in I} r \rightarrow s_i : \lambda_i ; \mathcal{H}_i} \quad \frac{\mathcal{H}\{\mu X. \mathcal{H}/X\} \xrightarrow{\alpha} \mathcal{G}}{\mu X. \mathcal{H} \xrightarrow{\alpha} \mathcal{G}}$$

Fig. 9. A transition relation between global types

induction may be applied, so $\mathcal{H}\{\mu X. \mathcal{H}/X\} \xrightarrow{p \rightarrow q : \lambda} \mathcal{G}'$ and $\mathbb{M} \vdash^g \mathcal{G}'$. By the third rule of Figure 9, $\mathcal{G} \xrightarrow{p \rightarrow q : \lambda} \mathcal{G}'$.

The remaining case is that $\mathcal{G} = \prod_{i \in I} r \rightarrow s_i : \lambda_i ; \mathcal{G}_i$. By unfolding the rules for transitions, we have

- $\text{PROC}(p, \mathbb{N}) = \ulcorner q! \lambda ; \mathbb{T}$.
- $\text{PROC}(q, \mathbb{N}) = \sum_{j \in J} p_j ? \lambda_j ; \mathbb{U}_j$, where $p = p_h$ and $\lambda = \lambda_h$ for some $h \in J$.

Furthermore, $\text{PROC}(p, \mathbb{M}) = \mathbb{T}$, $\text{PROC}(q, \mathbb{M}) = \mathbb{U}_h$ and $\text{PROC}(u, \mathbb{M}) = \text{PROC}(u, \mathbb{N})$ otherwise.

First assume that $p = r$. Then $\mathcal{G} \upharpoonright_p = \bigoplus_{i \in I} s_i ! \lambda_i ; (\mathcal{G}_i \upharpoonright_p)$. Hence, by the type rule for \bigoplus , since $\ulcorner q! \lambda ; \mathbb{T} \vdash \mathcal{G} \upharpoonright_p$, there is a $k \in I$ with $s_k = q$, $\lambda_k = \lambda$ and $\mathbb{T} \vdash \mathcal{G}_k \upharpoonright_p$. We have that $\mathcal{G} \xrightarrow{p \rightarrow q : \lambda} \mathcal{G}_k$. It remains to show that $\mathbb{M} \vdash^g \mathcal{G}_k$.

Since \mathcal{G} is closed, so is \mathcal{G}_k . Moreover, $\text{parties}(\mathcal{G}_k) \subseteq \text{parties}(\mathcal{G}) \subseteq \text{loc}(\mathbb{N}) = \text{loc}(\mathbb{M})$. Thus it remains to show that for each $u \in \text{loc}(\mathbb{M})$ one has $\text{PROC}(u, \mathbb{M}) \vdash \mathcal{G}_k \upharpoonright_u$ and $\mathcal{G}_k \upharpoonright_u$ is guarded. When $u = p$, we have $\text{PROC}(p, \mathbb{M}) = \mathbb{T} \vdash \mathcal{G}_k \upharpoonright_p$.

When $u \neq p, q$, we have $\mathcal{G} \upharpoonright_u = \prod_{i \in I} (p \rightarrow s_i : \lambda_i ; \mathcal{G}_i) \upharpoonright_u$. Since $\text{PROC}(u, \mathbb{M}) = \text{PROC}(u, \mathbb{N}) \vdash \mathcal{G} \upharpoonright_u$, by the rule for the merge in Figure 4, $\text{PROC}(u, \mathbb{M}) \vdash (p \rightarrow q : \lambda ; \mathcal{G}_k) \upharpoonright_u = \mathcal{G}_k \upharpoonright_u$.

Similarly, when $u = q$, $\text{PROC}(q, \mathbb{N}) \vdash (p \rightarrow q : \lambda ; \mathcal{G}_k) \upharpoonright_q = p ? \lambda ; \mathcal{G}_k \upharpoonright_q$. As $\text{PROC}(q, \mathbb{N}) = \sum_{j \in J} p_j ? \lambda_j ; \mathbb{U}_j$, there must be an $l \in J$ with $p_l = p$, $\lambda_l = \lambda$ and $\mathbb{U}_l \vdash \mathcal{G}_k \upharpoonright_q$. Now $\mathbb{N} \xrightarrow{p \rightarrow q : \lambda} \mathbb{M}'$, where $\text{PROC}(q, \mathbb{M}') = \mathbb{U}_l$. Since \mathbb{N} is race-free, $\mathbb{M}' = \mathbb{M}$ and thus $\mathbb{U}_l = \mathbb{U}_h = \text{PROC}(q, \mathbb{M})$. Hence $\text{PROC}(q, \mathbb{M}) \vdash \mathcal{G}_k \upharpoonright_q$.

In all these cases $\mathcal{G}_k \upharpoonright_u$ is a simple subterm of $\mathcal{G} \upharpoonright_u$, not within a recursion construct, so $\mathcal{G}_k \upharpoonright_u$ is guarded because $\mathcal{G} \upharpoonright_u$ is guarded.

Next assume that $p \neq r$. Observe that, for all $i \in I$, we have $s_i \neq p$; otherwise the first actions in $\text{PROC}(p, \mathbb{N})$ would be an external choice of receive actions, which is impossible.

The thread $\text{PROC}(r, \mathbb{N})$, possibly after unfolding recursion, must have the form $\bigoplus_{i \in I_0} s_i ! \lambda_i ; \mathbb{T}_i$ with $\mathbb{T}_i \vdash \mathcal{G}_i \upharpoonright_r$; here $I_0 \subseteq I$. For each $i \in I_0$ we define a network state \mathbb{N}_i such that $\mathbb{N}_i \vdash^g \mathcal{G}_i$. Take $\text{PROC}(r, \mathbb{N}_i) := \mathbb{T}_i$. The thread $\text{PROC}(s_i, \mathbb{N})$, possibly after unfolding recursion, must be of the form $r ? \lambda_i ; \mathbb{V}_i + \mathbb{U}_i$, where $\mathbb{V}_i \vdash \mathcal{G}_i \upharpoonright_{s_i}$; we take $\text{PROC}(s_i, \mathbb{N}_i) := \mathbb{V}_i$. For $u \neq r, s_i$ take $\text{PROC}(u, \mathbb{N}_i) := \text{PROC}(u, \mathbb{N})$. Since $\mathcal{G} \upharpoonright_u = \prod_{i \in I} \mathcal{G}_i \upharpoonright_u$ and $\text{PROC}(u, \mathbb{N}_i) \vdash \mathcal{G} \upharpoonright_u$, we have $\text{PROC}(u, \mathbb{N}_i) \vdash \mathcal{G}_i \upharpoonright_u$. Note that \mathcal{G}_i is closed since \mathcal{G} is closed, and $\text{parties}(\mathcal{G}_i) \subseteq \text{parties}(\mathcal{G}) \subseteq \text{loc}(\mathbb{N}) = \text{loc}(\mathbb{N}_i)$. Moreover, for $u \in \text{loc}(\mathbb{N})$, $\mathcal{G}_i \upharpoonright_u$ is guarded since $\mathcal{G} \upharpoonright_u$ is guarded. It follows that indeed $\mathbb{N}_i \vdash^g \mathcal{G}_i$.

For each $i \in I_0$ we have $\mathbb{N} \xrightarrow{\tau} * \xrightarrow{r \rightarrow s_i : \lambda_i} \mathbb{N}_i$. Hence \mathbb{N}_i is race-free. As $p, q \notin \{r, s_i \mid i \in I_0\}$, by race-freeness of \mathbb{N} , it follows that $\mathbb{N}_i \xrightarrow{p \rightarrow q : \lambda} \mathbb{M}_i$, where $\text{PROC}(p, \mathbb{M}_i) = \mathbb{T}$, $\text{PROC}(q, \mathbb{M}_i) = \mathbb{U}_h$, $\text{PROC}(r, \mathbb{M}_i) = \mathbb{T}_i$, $\text{PROC}(s_i, \mathbb{M}_i) = \mathbb{V}_i$, and $\text{PROC}(u, \mathbb{M}_i) = \text{PROC}(u, \mathbb{N})$ for all $u \notin \{p, q, r, s_i\}$. Furthermore $\|\mathcal{G}_i\|_p < \|\mathcal{G}\|_p$.

By the induction hypothesis there are \mathcal{G}'_i , for $i \in I_0$, with $\mathcal{G}_i \xrightarrow{p \rightarrow q : \lambda} \mathcal{G}'_i$ and $\mathbb{M}_i \vdash^g \mathcal{G}'_i$. Thus, by the second rule of Figure 9, $\mathcal{G} \xrightarrow{p \rightarrow q : \lambda} \mathcal{G}' := \prod_{i \in I_0} r \rightarrow s_i : \lambda_i ; \mathcal{G}'_i$. Trivially, \mathcal{G}' is closed and $\text{parties}(\mathcal{G}') \subseteq \text{loc}(\mathbb{M})$. Moreover, $\mathcal{G}' \upharpoonright_u$ is guarded for all $u \in \text{loc}(\mathbb{N})$. It remains to show that $\text{PROC}(u, \mathbb{M}) \vdash \mathcal{G}' \upharpoonright_u$ for all $u \in \text{loc}(\mathbb{M})$.

We have $\mathcal{G}' \upharpoonright_p = \prod_{i \in I_0} \mathcal{G}'_i \upharpoonright_p$. Since $\mathbb{M}_i \vdash^g \mathcal{G}'_i$, we have $\text{PROC}(p, \mathbb{M}_i) = \mathbb{T} \vdash \mathcal{G}'_i \upharpoonright_p$ for all $i \in I_0$. Thus, by the typing rule for merge, $\text{PROC}(p, \mathbb{M}) = \mathbb{T} \vdash \mathcal{G}' \upharpoonright_p$.

Likewise, $\mathcal{G}' \upharpoonright_q = \prod_{i \in I_0} \mathcal{G}'_i \upharpoonright_q$, $\text{PROC}(q, \mathbb{M}_i) = \mathbb{U}_h \vdash \mathcal{G}'_i \upharpoonright_q$ for all $i \in I_0$, and $\text{PROC}(q, \mathbb{M}) = \mathbb{U}_h \vdash \mathcal{G}' \upharpoonright_q$.

We have $\mathcal{G}' \upharpoonright_r = \bigoplus_{i \in I_0} s_i ! \lambda_i ; (\mathcal{G}'_i \upharpoonright_r)$. Since $\mathbb{M}_i \vdash^g \mathcal{G}'_i$, we have $\text{PROC}(r, \mathbb{M}_i) = \mathbb{T}_i \vdash \mathcal{G}'_i \upharpoonright_r$ for all $i \in I_0$. Thus, by the typing rule for internal choice, $\bigoplus_{i \in I_0} s_i ! \lambda_i ; \mathbb{T}_i \vdash \mathcal{G}' \upharpoonright_r$. Thus $\text{PROC}(r, \mathbb{M}) = \text{PROC}(r, \mathbb{N}) \vdash \mathcal{G}' \upharpoonright_r$.

For $u \neq p, q, r$ we have $\mathcal{G}' \upharpoonright_u = \prod_{i \in I_0} (r \rightarrow s_i : \lambda_i ; \mathcal{G}'_i) \upharpoonright_u$. Hence we need to show that $\text{PROC}(u, \mathbb{M}) \vdash (r \rightarrow s_i : \lambda_i ; \mathcal{G}'_i) \upharpoonright_u$ for all $i \in I_0$. So, pick $i \in I_0 \subseteq I$.

First suppose $u \neq s_i$. Since $\text{PROC}(u, \mathbb{M}) = \text{PROC}(u, \mathbb{N}) = \text{PROC}(u, \mathbb{M}_i)$ and $\mathbb{M}_i \vdash^g \mathcal{G}'_i$, we have $\text{PROC}(u, \mathbb{M}) \vdash \mathcal{G}'_i \upharpoonright_u = (r \rightarrow s_i : \lambda_i ; \mathcal{G}'_i) \upharpoonright_u$.

Finally, suppose $u = s_i$. As $\mathbb{M}_i \vdash^g \mathcal{G}'_i$, we have $\text{PROC}(s_i, \mathbb{M}_i) = \mathbb{V}_i \vdash \mathcal{G}'_i \upharpoonright_{s_i}$. By the typing rule for external choice, $r ? \lambda_i ; \mathbb{V}_i + \mathbb{U}_i \vdash r ? \lambda_i ; (\mathcal{G}'_i \upharpoonright_{s_i}) = (r \rightarrow s_i : \lambda_i ; \mathcal{G}'_i) \upharpoonright_{s_i}$. Hence $\text{PROC}(s_i, \mathbb{M}) = \text{PROC}(s_i, \mathbb{N}) \vdash (r \rightarrow s_i : \lambda_i ; \mathcal{G}'_i) \upharpoonright_{s_i}$. \square

Observation 4: The second rule in Figure 9 allows the index set to be narrowed. To understand why, consider the following global type.

$$\mathcal{G} \triangleq \begin{array}{l} (r \rightarrow t : a ; p \rightarrow q : a ; q \rightarrow r : a ; r \rightarrow s : a ; s \rightarrow q : a) \\ \boxplus r \rightarrow s : a ; s \rightarrow q : a ; p \rightarrow q : a ; q \rightarrow r : a ; r \rightarrow t : a \end{array}$$

Global type \mathcal{G} guardedly types the following network \mathbb{N} .

$$\begin{array}{l} p \llbracket q!a \rrbracket \\ \parallel q \llbracket (p?a; r!a; s?a) + (s?a; p?a; r!a) \rrbracket \\ \parallel r \llbracket t!a; q?a; s!a \rrbracket \\ \parallel s \llbracket r?a; q!a \rrbracket \\ \parallel t \llbracket r?a \rrbracket \end{array}$$

Network \mathbb{N} is race-free and $\mathbb{N} \xrightarrow{p \rightarrow q : a} \mathbb{M}$. If we insisted that $I = J$ in Fig. 9, then there would be no \mathcal{G}' such that $\mathcal{G} \xrightarrow{p \rightarrow q : a} \mathcal{G}'$ and $\mathcal{G}' \vdash^g \mathbb{M}$, as required for session fidelity. Narrowing of the global type by hiding branches of a choice, as permitted by $I \subseteq J$, is required when we have a race-free network, but the global type is not race-free, as in the above example.

Using the above, we can prove our soundness result. This is where we appeal to justness.

Theorem 3: If \mathbb{N} is guardedly well-typed and race-free, then $\mathbb{N} \models \mathcal{L}(J)$.

Proof: Let $\mathbb{N} = \mathbb{N}_0$ be race-free and assume that $\mathbb{N}_0 \vdash^g \mathcal{G}_0$. Let $\pi = \mathbb{N}_0 \xrightarrow{\tau, * p_0 \rightarrow q_0 : \lambda_0} \mathbb{N}_1 \xrightarrow{\tau, * p_1 \rightarrow q_1 : \lambda_1} \dots$ be a path on which some location $p \in \text{loc}(\mathbb{N}) = \text{loc}(\mathbb{N}_i)$ does not successfully terminate, and that contains only finitely many transitions involving p . We aim to show that π is not just.

Let $\ell(\pi) \in \mathbb{N}$ be the index n of the last state \mathbb{N}_n in this path, or $\ell(\pi) = \infty$ if π is infinite. By Lemmas 1, 2 and 8 there is a sequence $\mathcal{G}_0 \xrightarrow{p_0 \rightarrow q_0 : \lambda_0} \mathcal{G}_1 \xrightarrow{p_1 \rightarrow q_1 : \lambda_1} \dots$ of length $\ell(\pi)$ such that $\mathbb{N}_i \vdash^g \mathcal{G}_i$ for all i .

First consider the special case that for some \mathcal{G}_k in this sequence we have $p \notin \text{parties}(\mathcal{G}_k)$. Since \mathcal{G}_k is closed and $\text{PROC}(p, \mathbb{N}_k) \vdash \mathcal{G}_k \upharpoonright_p$, Lemma 7 yields that $\text{PROC}(p, \mathbb{N}_k) \vdash \text{OK}$. This implies that $\text{PROC}(p, \mathbb{N}_k)$ must have the form OK or $\mu X.\text{OK}$. As we assumed that p does not successfully terminate on π , it must stay a τ -transition away from successful termination. As this τ -transition is local to p , it follows that π is not just.

Thus we may assume that $p \in \text{parties}(\mathcal{G}_k)$ for all \mathcal{G}_k in the above sequence. By Corollary 3, $\|\mathcal{G}_i\|_p$ is finite for all i .

Claim: When $\mathbb{N} \vdash^g \mathcal{G}$, $\mathcal{G} \xrightarrow{t \rightarrow q : \lambda} \mathcal{H}$, $p \in \text{parties}(\mathcal{G})$ and $p \neq t, q$ then $\|\mathcal{H}\|_p \leq \|\mathcal{G}\|_p$. Moreover, if the transition $\mathcal{G} \xrightarrow{t \rightarrow q : \lambda} \mathcal{H}$ is derived without using the second rule in Figure 9, then $\|\mathcal{H}\|_p < \|\mathcal{G}\|_p$.

Proof: A trivial induction on the derivation of $\mathcal{G} \xrightarrow{t \rightarrow q : \lambda} \mathcal{H}$. Note that the conclusion $\|\mathcal{H}\|_p < \|\mathcal{G}\|_p$ is not warranted when

the second rule is used, due to the possibility that $p = r$ or $p = s_i$, where r and s_i are location variables of that rule. ■

Application of the claim: Since we assumed that π contains only finitely many transitions involving p , by restricting attention to a suffix of π we may just as well assume that no transition in π involves p , i.e., all p_i and q_i differ from p .

As $\|\mathcal{G}_i\|_p \geq 0$ for all i , there must be a \mathcal{G}_k in the above sequence such that $\|\mathcal{G}_l\|_p = \|\mathcal{G}_k\|_p$ for all $k \leq l \leq \ell(\pi)$, with $l \in \mathbb{N}$. So, past \mathcal{G}_k , all transitions $\mathcal{G}_l \xrightarrow{p_l \rightarrow q_l : \lambda_l} \mathcal{G}_{l+1}$ are derived by means of the second rule of Figure 9. Note that $\mathcal{G}_k \neq \text{OK}$ since $p \in \text{parties}(\mathcal{G})$, and $\mathcal{G}_k \neq X$ since \mathcal{G}_k is closed. Thus, possibly after unfolding recursion, \mathcal{G}_k must have the form $\boxplus_{j \in J} r \rightarrow s_j : \mu_j ; \mathcal{H}_j$. Since $\mathbb{N}_k \vdash^g \mathcal{G}_k$, we have $\text{PROC}(r, \mathbb{N}_k) \vdash \bigoplus_{j \in J} s_j ! \mu_j ; (\mathcal{H}_j \upharpoonright_r)$ and $\text{PROC}(s_j, \mathbb{N}) \vdash r ? \mu_j ; (\mathcal{H}_j \upharpoonright_{s_j})$ for each $j \in J$. In case $\text{PROC}(r, \mathbb{N}_k)$ never performs the τ -transitions needed to reach a thread state $\lceil s_j ! \mu_j ; T_j$ with $j \in J$, the path π is not just, and we are done. Likewise s_j will reach a state where it is ready to receive μ_j from r . So, for some $k \leq l \leq \ell(\pi)$, with $l \in \mathbb{N}$ and $j \in J$, we have $\mathbb{N}_l \xrightarrow{r \rightarrow s_j : \mu_j} \dots$.

A straightforward induction on $l \leq m < \ell(\pi)$ shows that \mathcal{G}_m has the form $\boxplus_{j \in J_m} r \rightarrow s_j : \mu_j ; \mathcal{H}_j^m$ with $j \in J_m \subseteq J$, so that $p_m \neq r$ and $q_m \neq s_j$. Here $j \in J_m$ follows since $\mathbb{N}_m \vdash^g \mathcal{G}_m$ and thus $\lceil s_j ! \mu_j ; T_j \vdash \mathcal{G}_m \upharpoonright_r$, and $p_m, q_m \neq r, s_j$ follows from the side condition in the second rule of Figure 9. It follows that in the path π , past \mathbb{N}_l neither location r nor s_j makes progress, and the transition $r \rightarrow s_j : \mu_j$ remains enabled. Hence π is not just. □