

Kent Academic Repository

Full text document (pdf)

Citation for published version

Bocchi, Laura and Demangeon, Romain (2013) Embedding Session Types in HML. In: Programming Language Approaches to Concurrency and Communication-cEntric Software (PLACES), 23rd March, 2013, Rome, Italy.

DOI

<https://doi.org/10.4204/EPTCS.137.5>

Link to record in KAR

<http://kar.kent.ac.uk/43739/>

Document Version

UNSPECIFIED

Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

Enquiries

For any further enquiries regarding the licence status of this document, please contact:

researchsupport@kent.ac.uk

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

Embedding Session Types in HML

Romain Demangeon¹ and Laura Bocchi^{2*}

Queen Mary, University of London
 romaind@eecs.qmul.ac.uk
 University of Leicester
 lb148@mcs.le.ac.uk

Abstract. Recent work on the enhancement of multiparty session types with logical annotations enable the effective verification of properties on (1) the structure of the conversations, (2) the sorts of the messages, and (3) the actual values exchanged. In [2] we extend this work to enable the specification and verification of mutual effects of multiple cross-session interactions. Here we give a sound and complete embedding into the Hennessy-Milner logic to justify the expressiveness of the approach in [2] and to provide it with a logical background that will enable us to compare it with similar approaches.

1 Introduction

The Hennessy-Milner Logic (HML) is an expressive modal logic with a strong semantic characterisation [6] that enables the specification of arbitrary behavioural properties of processes. Recent work on the enhancement of multiparty session types with logical annotations [2, 3] addressed key challenges for logical specifications of processes, which were unexplored in the context of HML, such as the tractability of specifications of multiparty choreographies.

The work in [2, 3] is based on multiparty session types [3, 4, 7] and inherits the same top-down approach. The key idea is that conversations are built as the composition of units of design called *sessions* which are specified from a global perspective (i.e., a global type). Each global type is then *projected*, making the responsibilities of each endpoint explicit. This approach enables: (1) the effective verification of properties such as session fidelity, progress, and error freedom, and (2) the modular local verification (i.e., of each principal) of global properties of multiparty interactions.

The direct use of HML for the same purpose would require to start from endpoint specifications and then to check their mutual consistency, hence would not offer the same tractability. Starting from global assertions, instead, results in significant concision, while still enjoying generality in the modelling and verification of choreographies.

By enhancing multiparty session types with logical annotations, [3] enables the effective verification of properties on the actual values exchanged, other than the properties on the sorts of the messages guaranteed by [4, 7]. For instance, global type G in (1) describes, following a similar syntax to [4], a conversation where role S sends role C an integer and then continues as specified by global type G' . Following [3], assertion \mathcal{G} in (1) can be obtained by annotating global type G ; assertion \mathcal{G} further prescribes that the

*This work has been partially sponsored by the project Leverhulme Trust award "Tracing Networks".

exchanged value, say y , must be greater than 10. Note that y is bound in \mathcal{G}' and the fact $\{y > 10\}$ can be relied on in the subsequent interactions.

$$G = S \rightarrow C : (\text{int}).G' \quad \mathcal{G} = S \rightarrow C : (y : \text{int})\{y > 10\}.G' \quad (1)$$

In [2] we extended [3] with the capability to refer to *virtual states* local to each network principal, hence expressing not only properties confined to the single multi-party sessions, but also stateful specifications incorporating mutual effects of multiple sessions run by a principal.

$$S \rightarrow C : (y : \text{int})\{y > 10 \wedge y = S.x\}\langle S.x++ \rangle \quad (2)$$

Consider now the protocol in (2). The description of this simple distributed application implies behavioural constraints of greater depth than the basic communication actions. The (sender-side) *predicate and effect* for the interaction step, $\{y > 10 \wedge y = S.x\}\langle S.x++ \rangle$, asserts that the message y sent to each client must equal the current value of $S.x$, a state variable x allocated to the *principal* serving as S ; and that the local effect of this message send is to increment $S.x$. In this way, S is specified to send incremental values across *consecutive* sessions. The resulting global specifications are called *multi-party stateful assertions* (MPSAs), and model the skeletal structure of the interactions of a session, the constraints on the exchanged messages and on the branches to be followed, and the *effects* of each interaction on the virtual state.

In order to obtain a clear understanding of the status of the logical methodology proposed in [2], it is useful to relate its notion of assertions to a more standard approach in process logic. This would enable us to integrate different methods catering for different concerns, for which we may need a common logical basis. In this paper we consider HML with predicates in [1, 3], and we justify the relevance of the stateful logical layer of [2] by embedding the behaviours of each role in a session – i.e., the projections of MPSAs – into a HML formula. In this way, the required predicates will hold if a process and its state perform reductions and updates matching those of the specification.

$$\forall y : \text{Nat}, [s_c(y)](y = S.x \wedge [S.x++]\text{true}) \quad (3)$$

(3) is the formula corresponding to the behaviour of S in (2) on channel s , where $[\ell]\phi$ means “if a process and its state perform the action ℓ , the resulting pair satisfies ϕ ”. Communications and state updates are treated as actions of a labelled transition system.

We explain how specifications handling several roles in several sessions can be soundly and completely embedded, through the use of an *interleaving* of formulae, exploring all the possible orders in which the actions coming from different sessions can be performed, and ensuring that predicates are always satisfied.

2 HML Embedding

MPSAs. We focus here on local assertions, ranged over by \mathcal{L} . \mathcal{L} refers to a specific role; we assume it derives, via projection, from some global assertion as in [2] – e.g., as (2).

$$\begin{aligned} \mathcal{L} ::= & \mathbf{p}\{l_i(x_i : U_i)\{A_i\}\langle E_i \rangle.\mathcal{L}_i\}_{i \in I} \mid \mathbf{p}^?\{l_i(x_i : U_i)\{A_i\}\langle E_i \rangle.\mathcal{L}_i\}_{i \in I} \\ & \mu t\{y : A'\}(x : S).\mathcal{L} : A \mid t(y : A') \mid \mathbf{end} \end{aligned}$$

Selection $\mathbf{p}!\{l_i(x_i : U_i)\{A_i\}\langle E_i \rangle.\mathcal{L}_i\}_{i \in I}$ models an interaction where the role sends \mathbf{p} a branch label l_i and a message x_i of sort U_i (e.g., `int`, `bool`, etc., and local assertion for delegation). A_i are predicates¹ and E_i state updates. We use guarded recursion, defining a parameter x initially set equal to a value satisfying the initialisation predicate A' , with A being an invariant predicate. Judgements are of the form $\Gamma \vdash P : \Delta$ where Γ is a shared environment and Δ is the session environment (see [2] for more details).

HML. [2] uses local assertions as a basis for the verification of a processes, ranged over by P . Here, the behaviour prescribed for P is modelled using the standard HML with the first-order predicates as in [1]. We use the same type of predicate A as in MP-SAs. We associate this HML with a LTS where actions ℓ model communications and state updates. Namely, $P, \sigma \xrightarrow{\ell} P', \sigma'$ if either $P \xrightarrow{\ell} P'$ and $\sigma' = \sigma$ or $P = P'$ and $\sigma' = \sigma$ after ℓ . We use ϕ to denote HML-formulae, which are built from predicates, implications, universal quantifiers, conjunctions and *must* modalities. We remark that the logic used in this *safety embedding* is positive: if we remove the implication symbol, there is no negation, no existential quantifier, no disjunction and no may modality. Additionally, the implication will always appear as $A \Rightarrow \phi$ meaning that modalities never appear in the negative side.

$$\phi ::= \mathbf{true} \mid \phi \wedge \phi \mid \phi \Rightarrow \phi \mid [\ell]\phi \mid A \mid \forall x : S.\phi \quad \ell ::= s[\mathbf{p}, \mathbf{q}](x) \mid \overline{s[\mathbf{p}, \mathbf{q}]}(x) \mid E$$

The satisfactions rules (Figure 1) are fairly standard, for a pair P, σ to satisfy a predicate A , A has to hold w.r.t. to σ , denoted by $\sigma \vdash_{\text{bool}} A$, meaning that $\sigma(A)$ is a tautology for the boolean logic.

$$\frac{\frac{P, \sigma \models \phi_1 \quad P, \sigma \models \phi_2}{P, \sigma \models \phi_1 \wedge \phi_2} \quad \frac{}{P, \sigma \models \mathbf{true}} \quad \frac{\text{if } P, \sigma \models \phi_1 \text{ then } P, \sigma \models \phi_2}{P, \sigma \models \phi_1 \Rightarrow \phi_2}}{\text{For all } P', \sigma' \text{ s.t. } P, \sigma \xrightarrow{\ell} P', \sigma', P', \sigma' \models \phi \quad \sigma \vdash_{\text{bool}} A \quad \text{For all values } v \text{ of type } T, P, \sigma \models \phi[v/x]} \frac{}{P \models [\ell]\phi} \quad \frac{}{P, \sigma \models A} \quad \frac{}{P, \sigma \models \forall x : T.\phi}$$

Fig. 1. Logical rules

The embedding of local types we propose is parametrised with a session channel $s[\mathbf{p}]$. Assertions appearing in input prefixes are embedded as premises in implications, and assertions in output prefixes have to be satisfied, yielding:

$$\|\mathbf{q}!\{l_i(x_i : S_i)\{A_i\}\langle E_i \rangle.\mathcal{L}_i\}_{i \in I}\|^{s[\mathbf{p}]} = \bigwedge_{i \in I} \forall x_i : S_i, [s[\mathbf{p}, \mathbf{q}]](x_i)(A_i \wedge [E_i]\|\mathcal{L}_i\|^{s[\mathbf{p}]}) \quad (4)$$

$$\|\mathbf{q}?\{l_j(x_j : S_j)\{A_j\}\langle E_j \rangle.\mathcal{L}_j\}_{j \in J}\|^{s[\mathbf{p}]} = \bigwedge_{j \in J} \forall x_j : S_j, [s[\mathbf{q}, \mathbf{p}]](x_j)(A_j \Rightarrow \|\mathcal{L}_j\|^{s[\mathbf{p}]}) \quad (5)$$

The embedding of selection (4), is a conjunction of formulae corresponding to the branches: for each value sent on the session channel, predicates should be satisfied and, if the state is updated, the embedding of the continuation should hold. For branching types (4), the assertion is used as an hypothesis and no update appear.

¹ As in [2, 3] we assume that the validity of closed formulae is decidable.

3 Soundness

To obtain soundness for typing judgements involving specifications, we have to introduce *interleavings* of formulae, treating the fact that one process can play several roles in several sessions. As a simple example both $s[\mathbf{p}_1, \mathbf{p}_2]^?(x).k![\mathbf{q}_1, \mathbf{q}_2] \langle 10 \rangle$ and $k![\mathbf{q}_1, \mathbf{q}_2] \langle 10 \rangle . s[\mathbf{p}_1, \mathbf{p}_2]^?(x)$ can be typed with $s[\mathbf{p}_2] : \mathbf{p}_1^?(x : \text{Nat}).\text{end}, k[\mathbf{q}_1] : \mathbf{q}_2!(y : \text{Nat}).\text{end}$.

Interleaving is not a new operator *per se* and can be seen as syntactic sugar, describing shuffling of must modalities. The main rule for interleaving is: $[\ell_1]\phi_1 \times [\ell_2]\phi_2 = [\ell_1](\phi_1 \times [\ell_2]\phi_2) \wedge [\ell_2](\phi_1 \wedge [\ell_1]\phi_2)$. When interleaving two or more formulae containing modalities, we obtain a conjunction of formulae, each one representing a different way of organising all modalities in a way preserving their initial orders. Informally, the interleaving of $[1][2]$ and $[A][B]$ is $[1][2][A][B] \wedge [A][B][1][2] \wedge [1][A][2][B] \wedge [A][1][B][2] \wedge [1][A][B][2] \wedge [A][1][2][B]$.

We encode a pair Δ, Γ into a complex formula $\text{Inter}(\Delta, \Gamma)$, defined as the interleaving of the formulae obtained by encoding the local types of Δ on their corresponding channels and the formulae corresponding to Γ , built as follows: for each channel $a : \mathbb{I}(\mathcal{G})$, if some $s[\mathbf{p}]$ is received on a , the resulting process should satisfy the encoding on $s[\mathbf{p}]$ of the projection of \mathcal{G} on \mathbf{p} :

$$\text{Inter}(s_1[\mathbf{p}_1], \dots, s_n[\mathbf{p}_n]; a_1 : \mathbb{I}(\mathcal{G}_1), \dots, a_m : \mathbb{I}(\mathcal{G}_m)) \\ = \|T_1\|^{s_1[\mathbf{p}_1]} \times \dots \times \|T_n\|^{s_n[\mathbf{p}_n]} \times \phi_1 \times \dots \times \phi_m$$

where $\phi_i = \forall s'_i. \forall \mathbf{p}'_i. [a_i(s'_i[\mathbf{p}_i])] \| \mathcal{G}_i \upharpoonright \mathbf{p}'_i \|^{s'_i[\mathbf{p}'_i]}$.

The erasing operator $\text{Er}(\mathcal{L})$, which translates an asserted type into its unasserted counterpart is defined by the removal of every assertion from the local types. Unasserted typing rules for the judgements $\vdash P \triangleright \Delta$ are easily deduced. Our preciseness result is:

Proposition 1 (Preciseness). *If $\Gamma \vdash P \triangleright \Delta$, then: $P, \sigma \models (\text{Inter}(\Delta, \Gamma))$. If $\vdash P \triangleright \text{Er}(\Delta)$ and $P, \sigma \models (\text{Inter}(\Delta, \Gamma))$ then $\Gamma \vdash P \triangleright \Delta$.*

4 Refinements

Embedding to pure HML We are actually able to embed a stateful satisfaction relation $P, \sigma \models \phi$ into a satisfaction relation $P' \models \phi'$ for a standard π -calculus with first-order values, by encoding the store σ into a π -process:

$$\|x_1 \mapsto v_1, \dots, x_n \mapsto v_n\|_{\mathbf{p}} = \overline{a_1}(v_1) \mid \dots \mid \overline{a_n}(v_n) \mid \\ !x_1(e).a_1(y_1) \dots a_n(y_n).(\overline{a_1}(\text{eval}(e[y_1 \dots y_n/x_1 \dots x_n]))) \mid \overline{a_2}(y_2) \mid \dots \mid \overline{a_n}(y_n) \mid \dots \mid \\ !x_n(e).a_1(y_1) \dots a_n(y_n).(\overline{a_1}(y_1) \mid \dots \mid \overline{a_{n-1}}(y_{n-1}) \mid \overline{a_n}(\text{eval}(e[y_1 \dots y_n/x_1 \dots x_n])))$$

For each variable x_i in the domain of the state σ , we add an output prefix emitting its content on the channel a_i and we add a replicated module that waits for an update e at x_i , then capture the value of all variables of the current state, replace the variable x_i by evaluating e by `eval`, and then makes available the other ones. Soundness and completeness allow us to state that HML formulae for pairs state/process can be seen as pure HML formulas on the π -processes.

The embedding for the formula is given by $\| [E] \phi \|_{\mathbf{p}} = \| [E] \|_{\mathbf{p}} \| \phi \|_{\mathbf{p}}$ and $\| A \|_{\mathbf{p}} = [\overline{a_1}(v_1)] \dots [\overline{a_n}(v_n)] A \{v_1, \dots, v_n/x_1, \dots, x_n\}$ where the state variables of A are x_1, \dots, x_n .

Proposition 2 (Preciseness). *If $P, \sigma \models \phi$, then $\|P\|_{\mathbb{P}} \mid \|\sigma\|_{\mathbb{P}} \models \|\phi\|_{\mathbb{P}}$.
If $\|P\|_{\mathbb{P}} \mid \|\sigma\|_{\mathbb{P}} \models \|\phi\|_{\mathbb{P}}$ then $P, \sigma \models \phi$*

Embedding Recursion Recursion, absent from the previous embeddings, can actually be encoded at the cost of much technical details. We add to our HML syntax the recursion operators, $\mu X.\phi$ and X (similar to the ones present in the μ -calculus [5]). The main difficulty lies in the interaction between interleaving and recursion: loops coming from different sessions can be interleaved in many different way, and the difficult task is to compute the finite formula which is equivalent to this interleaving. As a small example consider the following session environment (interactions are replaced by integer labels): $s_1[p_1] : \mu X.1.2.X, s_2[p_2] : \mu Y.3.4.Y$. The simplest HML formula describing all possible interleavings is:

$$\begin{aligned} & \mu A.([1]\mu B.([2]A \wedge [3]\mu C.([4]B \wedge [2]([1]C \wedge [4].A))) \wedge \\ & [3]\mu D.([4].A \wedge [1]\mu E.([2]D \wedge [4]([2]A \wedge [3]E)))) \end{aligned}$$

We will not detail here how we proceed. The idea is that we translate every session judgment into formula, and then every formula into a finite automaton where one state corresponds either to a sequence or a recursion, and where transitions correspond to inputs and outputs. The automata are then merged into a parallel automaton, which is expanded recursively into branch automata and translated back into formulas.

On our example, we obtain the formulas $\mu X.[1][2].X$ and $\mu Y.[3][4].Y$, each one giving an automaton with 2 states (initial and between [1] (resp. [3]) and [2] (resp. [4])). Merging yields automata with 4 states: the initial one, one after [1], one after [3], one after both [1] and [3]. This automata is diamond-shaped, and, as a result, not tree-shaped. Expansion yields an automata with 7 states, which is then translated in the formula described above. The preciseness proof relies on the fact that the operation described in 3. and 4. give equivalent automata, and that two formulas translated to two equivalent automata are equivalent for the HML satisfaction relation.

References

1. Martin Berger, Kohei Honda, and Nobuko Yoshida. Completeness and logical full abstraction in modal logics for typed mobile processes. In *ICALP (2)*, volume 5126 of *LNCS*, pages 99–111. Springer, 2008.
2. Laura Bocchi, Romain Demangeon, and Nobuko Yoshida. A multiparty multisession logic. (to appear in *TGC*), 2012.
3. Laura Bocchi, Kohei Honda, Emilio Tuosto, and Nobuko Yoshida. A theory of design-by-contract for distributed multiparty interactions. In *CONCUR*, volume 6269 of *LNCS*, pages 162–176, 2010.
4. Mario Coppo and Mariangiola Dezani-Ciancaglini. Structured communications with concurrent constraints. In *TGC*, pages 104–125, 2008.
5. Mads Dam. CTL* and ECTL* as fragments of the modal mu-calculus. *TCS*, 126(1):77–96, 1994.
6. Matthew Hennessy and Robin Milner. Algebraic Laws for Non-Determinism and Concurrency. *JACM*, 32(1), 1985.
7. Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty Asynchronous Session Types. In *POPL'08*, pages 273–284. ACM, 2008.