

HYPE: Hybrid modelling by composition of flows

Vashti Galpin¹, Luca Bortolussi² and Jane Hillston¹

¹Laboratory for Foundations of Computer Science, University of Edinburgh

²Department of Maths and Computer Science, University of Trieste

Abstract. Hybrid systems are manifest in both the natural and the engineered world, and their complex nature, mixing discrete control and continuous evolution, make it difficult to predict their behaviour. In recent years several process algebras for modelling hybrid systems have appeared in the literature, aimed at addressing this problem. These all assume that continuous variables in the system are modelled monolithically, often with differential equations embedded explicitly in the syntax of the process algebra expression. In HYPE an alternative approach is taken which offers finer-grained modelling with each flow or influence affecting a variable modelled separately. The overall behaviour then emerges as the composition of flows.

In this paper we give a detailed account of the HYPE process algebra, its semantics, and its use for verification of systems. We establish both syntactic conditions (well-definedness) and operational restrictions (well-behavedness) to ensure reasonable behaviour in HYPE models. Furthermore we consider how the equivalence relation defined for HYPE relates to other relations previously proposed in the literature, demonstrating that our fine-grained approach leads to a more discriminating notion of equivalence. We present the HYPE model of a standard hybrid system example, both establishing that our approach can reproduce the previously obtained results and demonstrating how our compositional approach supports variations of the problem in a straightforward and flexible way.

Keywords: hybrid systems; process algebra; flows; compositionality; bisimulation

1. Introduction

Predominantly, process algebras have been used to model discrete systems. Their style of modelling, in terms of agents and actions, matches well with a discrete event view of the world. This form of modelling encompasses a wide variety of systems as witnessed by the wide-range of applications which have emerged for process algebra modelling. Nevertheless there are situations in which this discrete view is not entirely satisfactory. For example there may be inherent properties of the system which vary continuously or because a fluid abstraction (in which discrete variables are treated as continuous) of some aspects of the system will bring benefits for modelling, or analysis, or both. We term a *hybrid system* one which exhibits both discrete

Correspondence and offprint requests to: Vashti Galpin, Informatics Forum, 10 Crichton Street, Edinburgh EH8 9AB. e-mail: Vashti.Galpin@ed.ac.uk

and continuous behaviour. Given the success and attractiveness of process algebra modelling it is appealing to investigate whether the same compositional approach can be applied to such hybrid systems.

Hybrid behaviour arises in a variety of systems, both engineered and natural. Consider a thermostatically controlled heater. The continuous variable is air temperature, and the discrete events are the switching on and off of the heater by the thermostat in response to the air temperature. Another example would be a genetic regulatory network, such as the Repressilator [EL00], in which genes can be switched on or off by interactions with their environment (more precisely, with transcription factor proteins). The behaviour of such systems can be regarded as a collection of sets of ODEs, the discrete events shifting the dynamic behaviour from the control of one set of ODEs to another and possibly modifying the values of the variables immediately after the event. This is the approach taken with hybrid automata [Hen96].

HYPE is a novel process algebra for modelling hybrid systems. A hybrid system consists of values which change continuously over time with respect to specific dynamics. Discrete events can cause discontinuous jumps in these values after which different dynamics may come into effect. These events can be triggered by conditions on the continuously changing values. When modelling a hybrid system we need to capture both the dynamics and the events which cause the system to move between different dynamics. The novelty of HYPE lies in how it captures the continuous dynamics of a system, its modular style of definition and the separation of a discrete controller considered in parallel to the system under study. Unlike previous process algebras for hybrid systems, HYPE captures behaviour at a fine-grained level, composing distinct flows or influences. The dynamic behaviour then emerges, via the semantics of the language, from these compositional elements. For the continuous aspects, we are inspired by the fluid flow semantics of PEPA models [Hil05, TGH10] which approximates the behaviour of large numbers of discrete components with a set of ordinary differential equations (ODEs). For the discrete aspects of the model, the controller imposes sequencing on events and having this separated readily supports investigations into how the same system may behave under the influence of distinct controllers. Note that although the controller is specified as a distinct component of the system to aid modularity of definition, its events are triggered by the continuous values of the system.

Process algebras have the advantage of being compositional hence models are built out of subcomponents and we fully exploit this feature in HYPE. Existing process algebras for hybrid systems include ACP_{hs}^{srt} [BM05], hybrid χ [vBMR⁺06], ϕ -calculus [RS03] and HyPA [CR05]. In [Kha06], Khadim shows substantial differences in the approaches taken by these process algebras relating to syntax, semantics, discontinuous behaviour, flow-determinism, theoretical results and availability of tools. However, they are all similar in their approach in that the dynamic behaviour of each subcomponent must be fully described with the ODEs for the subcomponent given explicitly in the syntax of the process algebra, before the model can be constructed. Furthermore an evaluation of continuous variables is inherent in the notion of state.

We aim for a finer-grained approach where each subcomponent is built up from a number of flows and hence the ODEs are only obtained once the model is constructed. By flow, we mean something that has an influence on a quantity of interest. For example, in a tank with two inlets and an outlet, both the inlets and the outlet influence the tank level, hence here we would identify three separate flows. The continuous part of the system is represented by the appropriate variables and the change over time of a given variable is determined by a number of active influences which represent flows and are additive in nature. Our approach also differs in that we explicitly require a controller that consists only of events.

We believe that the use of flows as the basic elements of model construction has advantages such as ease and simplification of modelling. This approach assists the modeller in allowing them to identify smaller or local descriptions of the model and then to combine these descriptions to obtain the larger system. The explicit controller also helps to separate modelling concerns.

In this paper we present a full account of the HYPE language, including both operational and hybrid semantics. The two are distinct because in HYPE the *state* of the system does not include an evaluation of the continuously varying variables; instead it captures the influences which are currently in operation within the system. Thus each state corresponds to one dynamic regime, captured by a set of ODEs in the hybrid semantics. We also establish an equivalence semantics for the language and consider how this compares with the existing bisimulations for hybrid process algebras and hybrid automata which have appeared in the literature. In particular we consider the bisimulations defined in [BM05] for ACP_{hs}^{srt} and the U -bisimulation of [AMP⁺03] for hybrid automata, showing that the notion of bisimulation in HYPE is more discriminating than that for hybrid automata.

This paper is a revised and extended version of the paper which appeared in [GBH09]. Here we have substantially revised the definition of what it means for a HYPE model to be well-defined and furthermore

considered when a model may be considered to be *well-behaved* and how to check for such a property. A model which is not well-behaved may exhibit *instantaneous Zeno behaviour*, completing an infinite number of events in a single time instant. The work on comparing bisimulations is extended and here, additionally, we compare the expressive power of HYPE with that of hybrid automata at the level of composition of models. Finally a new example of HYPE modelling that considers a train gate controller [AHH96], is presented.

The structure of the rest of the paper is as follows. In the next section we introduce our syntax for hybrid systems, explaining its components. In the following sections, we present the operational semantics and the hybrid semantics, explaining how we go from the notion of state to the ODEs which describe the system dynamics. We then identify the subclass of HYPE models which we consider to be well-defined and discuss how to check that a HYPE model is well-behaved. We also compare composition of models in HYPE and hybrid automata. We define a notion of bisimulation in Section 8 and compare it with previous bisimulations for hybrid languages from the literature. To illustrate the power of HYPE, in Section 10 we present an example, a train gate controller. The remaining sections discuss related work, conclusions and future research.

2. HYPE Definition

This section will present HYPE by way of a running example of the temperature control system of an orbiting spacecraft. As the orbiter travels around the earth, it needs to regulate its temperature to remain within operational limits. It has insulation but needs to use a heater at low temperatures and at high temperatures it can erect a shade to reduce temperatures. The modelling spirit of HYPE focusses on flows. In our example, we identify four flows affecting the temperature. One is due to thermodynamic cooling, one is due to the heater, one is due to the heating effect of the sun and one is due to the cooling effect of the shade. The strength and form of a flow are modified by events. We first define the heater which can be on or off.

$$Heat \stackrel{def}{=} \underline{on} : (h, r_h, const).Heat + \underline{off} : (h, 0, const).Heat + \underline{init} : (h, 0, const).Heat$$

This is a *summation* of prefixes. Each prefix consists of two actions. *Events* ($\underline{a} \in \mathcal{E}$) are actions which happen instantaneously and trigger discrete changes. They can be caused by a controller or happen randomly and can depend on the global state of the system, specifically values of variables. In the example, the events are on, switching on; off, switching off and init, the initialisation event.

Activities ($\alpha \in \mathcal{A}$) are *influences* on the evolution of the continuous part of the system and define flows. An activity is defined as a triple and can be parameterised by a set of variables, $\alpha(\mathcal{W}) = (\iota, r, I(\mathcal{W}))$. This triple consists of an *influence name* ι , a rate of change (or *influence strength*) r and an *influence type name* $I(\mathcal{W})$ which describes how that rate is to be applied to the variables involved, or the actual form of the flow¹. In *Heat*, there are two distinct activities, $(h, r_h, const)$ and $(h, 0, const)$. The first one gives the effect of the heater being on: the influence name is h which represents the influence from the heater on the orbiter's temperature, r_h is the strength of the heater, and it is associated with the function called *const*. The second captures the effect of the heater being off. It again affects influence h , it has strength 0 and the form it takes is *const*.

The interpretation of influence types will be specified separately, so that experimentation with different functional forms of the heating flow can occur without modifying the subcomponent. Hence, in HYPE we separate the description of the logical structure of flows from their mathematical interpretation. We now describe the other flows for the example.

$$\begin{aligned} Shade &\stackrel{def}{=} \underline{up} : (d, -r_d, const).Shade + \underline{down} : (d, 0, const).Shade + \underline{init} : (d, 0, const).Shade \\ Sun &\stackrel{def}{=} \underline{light} : (s, r_s, const).Sun + \underline{dark} : (s, 0, const).Sun + \underline{init} : (s, 0, const).Sun \\ Cool(K) &\stackrel{def}{=} \underline{init} : (c, -1, linear(K)).Cool(K) \end{aligned}$$

The only event in the last definition is the initialisation event, as once cooling is in effect it does not change. The influence name is c and its strength is -1 . The influence type *linear*(K) will be interpreted as a linear function of its variable K which represents the temperature of the system. We also need to model the change in sunlight. We do this by keeping track of time with the following component (which is kept simple for

¹ For convenience, we will use I for $I(\mathcal{W})$ when \mathcal{W} can be inferred.

reasons of space).

$$Time \stackrel{def}{=} \underline{\text{light}} : (t, 1, \text{const}). Time + \underline{\text{dark}} : (t, 1, \text{const}). Time + \underline{\text{init}} : (t, 1, \text{const}). Time$$

These subcomponents can be combined to give the overall uncontrolled system².

$$Sys \stackrel{def}{=} (((Heat \underset{\{\text{init}\}}{\boxtimes} Shade) \underset{\{\text{init}\}}{\boxtimes} Sun) \underset{\{\text{init}\}}{\boxtimes} Cool(K)) \underset{\{\text{init}, \text{light}, \text{dark}\}}{\boxtimes} Time$$

Here \boxtimes represents parallel cooperation. L is the set of events over which synchronisation must occur. Events not in L can occur independently. Sys is called the *uncontrolled system* because all events are possible and no causal or temporal constraints have been imposed yet. For instance, we need to specify that the heater can only be switched off after it has been switched on. We now give controllers/sequencers for the heater, the shade and the effect of the sun. Note that these correspond to the starting states set by init events. For example, in *Heat*, init sets the heater to off (the influence strength is 0), therefore on is the next event to occur.

$$\begin{aligned} Con_h &\stackrel{def}{=} \underline{\text{on.off}}. Con_h \\ Con_d &\stackrel{def}{=} \underline{\text{up.down}}. Con_d \\ Con_s &\stackrel{def}{=} \underline{\text{light.dark}}. Con_s \\ Con &\stackrel{def}{=} Con_h \underset{\emptyset}{\boxtimes} Con_d \underset{\emptyset}{\boxtimes} Con_s \end{aligned}$$

Controllers only have event prefixes. Their behaviour is affected by the state of the system through event conditions which determine when events occur. The controlled system is constructed from cooperation of the controller and the uncontrolled system and the controller must be prefixed by the initialisation event init. For the example, the controlled system is described by

$$TempCtrl \stackrel{def}{=} Sys \underset{M}{\boxtimes} \underline{\text{init}}. Con \quad \text{with} \quad M = \{\underline{\text{init}}, \underline{\text{on}}, \underline{\text{off}}, \underline{\text{up}}, \underline{\text{down}}, \underline{\text{light}}, \underline{\text{dark}}\}.$$

This has defined the structure of our system but we require additional definitions to capture further details. We need to link each influence with an actual variable. This is done using the function iv . For the example, $iv(h) = iv(s) = iv(d) = iv(c) = K$ where K is the actual variable for the temperature of the orbiter, and $iv(t) = T$, the variable for time. Note that an influence can only be associated with one variable, in agreement with the interpretation of influences as flows. This does not mean that only one variable can be affected by an event. For another variable Y that was also affected by the heater being on (power consumption, say), we could define a subcomponent with a prefix on:(p, r, I) and set $iv(p) = Y$.

We define the influence types as $\llbracket \text{const} \rrbracket = 1$ and $\llbracket \text{linear}(X) \rrbracket = X$. The influence types are used to describe how influences are affected by variables in the system. The type *const* is used when there is no dependency and *linear*(X) is used when the influence depends linearly on the value of the variable X . Non-linear dependencies can also be defined through this mechanism. Examples of non-linear ODEs occur in the description of the motion of a pendulum and when mass action is used in biological modelling.

Finally, we define what triggers an event, and how it affects variables, with the function ec . Each event condition consists of an activation condition and a reset.

The activation condition is a positive boolean formula (one without negation) containing equalities and nonstrict inequalities on system variables or the symbol \perp . We choose this specific restriction since it guarantees that the set represented by such a formula is a closed set, being the finite union and intersection of closed sets. We require closed sets so that the first instant in which a guard becomes true is defined in a consistent way.

The variable reset is a conjunction of equality predicates of the form $V' = f(V)$ where V' denotes the new value that V will have after the reset, while V denotes the previous value. Resets of the form $V' = V$ can be left implicit, and the identity reset is denoted *true*. Event conditions are associated with events by the function ec . The example has the following event conditions.

² Note that in HYPE uncontrolled systems all possible controllable behaviour is exhibited, awaiting a suitable controller to instigate appropriate triggers. This is in contrast to, say, an embedded system, where the uncontrolled system would be aspects of behaviour outside the scope of the controller.

$$\begin{array}{ll}
 ec(\underline{\text{init}}) & = (true, (K' = k_0 \wedge T' = 0)) \\
 ec(\underline{\text{off}}) & = ((K \geq k_1), true) \\
 ec(\underline{\text{up}}) & = ((K \geq k_3), true) \\
 ec(\underline{\text{light}}) & = ((T = 12), true) \\
 ec(\underline{\text{on}}) & = (K \leq k_2, true) \\
 ec(\underline{\text{down}}) & = (K \leq k_4, true) \\
 ec(\underline{\text{dark}}) & = (T = 24, T' = 0)
 \end{array}$$

where the $k_i (1 \leq i \leq 4)$ are fixed temperature values. Most events are urgent – the event must occur as soon as its event condition is satisfied. For events that can happen randomly such as breakdowns, we introduce a special event condition \perp which means that the event can happen nondeterministically in time (sometimes termed lazy jumps)³. In the example, the $\underline{\text{init}}$ event has an associated event condition of $true$ and so this must happen immediately and $\underline{\text{light}}$ happens when 12 hours have passed. The event $\underline{\text{init}}$ has a reset that defines the values of the variables and $\underline{\text{on}}$ has a reset of $true$ meaning that no values are changed. Figure 1 presents all components of the HYPE model of the orbiter in one place. For a comparison of different process algebras for modelling hybrid systems (other than HYPE) see [Kha06, TCT01] and we discuss this further in Section 9.

In the preceding informal discussion, we have introduced the main constituents of a HYPE model including the combination of flow components with a controller component, variables, association between influences and variables, conditions that specify when events occur, and definitions for the influence type functions. To understand the dynamics of this system, we need to derive ODEs to describe how the variables change over time. To do this we present operational semantics that define the behaviour of our controlled system. Before that we present the formal definition of HYPE. In the rest of the paper, \mathcal{V} is a set or tuple of variables with $\mathcal{W} \subseteq \mathcal{V}$ denoting an arbitrary subset of \mathcal{V} .

Definition 2.1. A *controlled system* is constructed as follows.

- *Subcomponents* are defined by $C_s(\mathcal{W}) \stackrel{\text{def}}{=} S$, where C_s is the *subcomponent name* and S satisfies the grammar $S' ::= \underline{a} : \alpha.C_s(\mathcal{W}) \mid S' + S' \ (\underline{a} \in \mathcal{E}, \alpha \in \mathcal{A})$, with the free variables of S in \mathcal{W} .
- *Components* are defined by $C(\mathcal{W}) \stackrel{\text{def}}{=} P$, where C is the *component name* and P satisfies the grammar $P' ::= C_s(\mathcal{W}) \mid P' \underset{L}{\bowtie} P'$, with the free variables of P in \mathcal{W} and $L \subseteq \mathcal{E}$.
- An *uncontrolled system* Σ is defined according to the grammar $\Sigma' ::= C_s(\mathcal{W}) \mid C(\mathcal{W}) \mid \Sigma' \underset{L}{\bowtie} \Sigma'$, where $L \subseteq \mathcal{E}$ and \mathcal{W} is a set of system variables.
- *Controllers* only have events: $M ::= \underline{a}.M \mid 0 \mid M + M$ with $\underline{a} \in \mathcal{E}$ and $Con ::= M \mid Con \underset{L}{\bowtie} Con$ with $L \subseteq \mathcal{E}$.
- A *controlled system* is $ConSys ::= \Sigma \underset{L}{\bowtie} \underline{\text{init}}.Con$ where $L \subseteq \mathcal{E}$. The set of controlled systems is \mathcal{C}_{Sys} .

A controlled system together with the appropriate sets and functions, gives a HYPE model.

Definition 2.2. A *HYPE model* is a tuple $(ConSys, \mathcal{V}, IN, IT, \mathcal{E}, \mathcal{A}, ec, iv, EC, ID)$ where

- $ConSys$ is a controlled system as defined above.
- \mathcal{V} is a finite set of variables⁴.
- IN is a set of influence names and IT is a set of influence type names.
- \mathcal{E} is a set of events of the form \underline{a} and \underline{a}_i .
- \mathcal{A} is a set of activities of the form $\alpha(\mathcal{W}) = (\iota, r, I(\mathcal{W})) \in (IN \times \mathbb{R} \times IT)$.
- $ec : \mathcal{E} \rightarrow EC$ maps events to event conditions. Event conditions are pairs of formulas, the first with free variables in \mathcal{V} and the second with free variables in $\mathcal{V} \cup \mathcal{V}'$.
- $iv : IN \rightarrow \mathcal{V}$ maps influence names to variable names.
- EC is a set of event conditions.
- ID is a collection of definitions consisting of a real-valued function for each influence type name $\llbracket I(\mathcal{W}) \rrbracket = f(\mathcal{W})$ where the variables in \mathcal{W} are from \mathcal{V} .
- $\mathcal{E}, \mathcal{A}, IN$ and IT are pairwise disjoint.

³ We have not done so here but it is possible to have events that are probabilistic rather than nondeterministic in time. Tuffin *et al* [TCT01] use a value drawn from an exponential distribution for the time until the next event and we consider the use of stochastic delays in [BGH10a].

⁴ In [GBH09], a set of formal variables \mathcal{X} was also defined but with experience, it became clear that these variables were not required and have been dropped.

$$\begin{aligned}
Heat &\stackrel{\text{def}}{=} \text{on} : (h, r_h, \text{const}).Heat + \text{off} : (h, 0, \text{const}).Heat + \text{init} : (h, 0, \text{const}).Heat \\
Shade &\stackrel{\text{def}}{=} \text{up} : (d, -r_d, \text{const}).Shade + \text{down} : (d, 0, \text{const}).Shade + \text{init} : (d, 0, \text{const}).Shade \\
Sun &\stackrel{\text{def}}{=} \text{light} : (s, r_s, \text{const}).Sun + \text{dark} : (s, 0, \text{const}).Sun + \text{init} : (s, 0, \text{const}).Sun \\
Cool(K) &\stackrel{\text{def}}{=} \text{init} : (c, -1, \text{linear}(K)).Cool(K) \\
Time &\stackrel{\text{def}}{=} \text{light} : (t, 1, \text{const}).Time + \text{dark} : (t, 1, \text{const}).Time + \text{init} : (t, 1, \text{const}).Time \\
\\
Sys &\stackrel{\text{def}}{=} (((Heat \boxtimes_{\{\text{init}\}} Shade) \boxtimes_{\{\text{init}\}} Sun) \boxtimes_{\{\text{init}\}} Cool(K)) \boxtimes_{\{\text{init}, \text{light}, \text{dark}\}} Time \\
\\
Con_h &\stackrel{\text{def}}{=} \text{on.off}.Con_h \\
Con_d &\stackrel{\text{def}}{=} \text{up.down}.Con_d \\
Con_s &\stackrel{\text{def}}{=} \text{light.dark}.Con_s \\
Con &\stackrel{\text{def}}{=} Con_h \boxtimes_{\emptyset} Con_d \boxtimes_{\emptyset} Con_s \\
\\
TempCtrl &\stackrel{\text{def}}{=} Sys \boxtimes_M \text{init}.Con \quad \text{with} \quad M = \{\text{init}, \text{on}, \text{off}, \text{up}, \text{down}, \text{light}, \text{dark}\} \\
\\
iv(t) &= T \qquad \qquad \qquad iv(h) = iv(s) = iv(d) = iv(c) = K \\
\\
ec(\text{init}) &= (\text{true}, (K' = k_0 \wedge T' = 0)) \\
ec(\text{off}) &= ((K \geq k_1), \text{true}) \qquad \qquad ec(\text{on}) = (K \leq k_2, \text{true}) \\
ec(\text{up}) &= ((K \geq k_3), \text{true}) \qquad \qquad ec(\text{down}) = (K \leq k_4, \text{true}) \\
ec(\text{light}) &= ((T = 12), \text{true}) \qquad \qquad ec(\text{dark}) = (T = 24, T' = 0) \\
\\
\llbracket \text{const} \rrbracket &= 1 \qquad \qquad \qquad \llbracket \text{linear}(X) \rrbracket = X
\end{aligned}$$

Fig. 1. HYPE model of the orbiter

In the sequel, we will use P, Q, \dots for controlled systems (rather than $ConSys$) and for HYPE models. When referring to a HYPE model, the meta-variable for the controlled system will be used, such as P , and the tuple will be implied. In the case of two HYPE models P and Q without reference to the tuple, we will assume two implied tuples with identical elements except for the first elements. The syntax of HYPE is moderately complex because hybrid systems are complex structures displaying both continuous and discrete behaviour. Despite the complexity, the example shows how it is straightforward to construct a HYPE model once flows and the controller are identified. Additionally the modularity of definition helps the modeller to check that everything is specified.

3. Operational Semantics

We now introduce the behaviour of HYPE models by defining an appropriate semantics. We start by considering the events that can happen, and construct a labelled transition system describing this behaviour. To define the operational semantics, a notion of state is required. States record for each influence, its current strength and influence type. Each configuration will have an associated state which will capture the continuous behaviour in that configuration.

Definition 3.1. An *operational state* of the system is a function $\sigma : IN \rightarrow (\mathbb{R} \times IT)$. The set of all operational states is \mathcal{S} . A *configuration* consists of a controlled system together with an operational state $\langle ConSys, \sigma \rangle$ and the set of configurations is \mathcal{F} .

We use ‘state’ for ‘operational state’ in the rest of this document. For convenience, states may be written as a set of triples of the form $(\iota, r, I(\mathcal{W}))$ where there is at most most one triple containing ι . This is the same

Prefix with influence:	$\langle \underline{a} : (\iota, r, I).E, \sigma \rangle \xrightarrow{\underline{a}} \langle E, \sigma[\iota \mapsto (r, I)] \rangle$
Prefix without influence:	$\langle \underline{a}.E, \sigma \rangle \xrightarrow{\underline{a}} \langle E, \sigma \rangle$
Choice:	$\frac{\langle E, \sigma \rangle \xrightarrow{\underline{a}} \langle E', \sigma' \rangle}{\langle E + F, \sigma \rangle \xrightarrow{\underline{a}} \langle E', \sigma' \rangle} \quad \frac{\langle F, \sigma \rangle \xrightarrow{\underline{a}} \langle F', \sigma' \rangle}{\langle E + F, \sigma \rangle \xrightarrow{\underline{a}} \langle F', \sigma' \rangle}$
Cooperation without synchronisation:	$\frac{\langle E, \sigma \rangle \xrightarrow{\underline{a}} \langle E', \sigma' \rangle}{\langle E \boxtimes_M F, \sigma \rangle \xrightarrow{\underline{a}} \langle E' \boxtimes_M F, \sigma' \rangle} (\underline{a} \notin M)$
Cooperation with synchronisation:	$\frac{\langle F, \sigma \rangle \xrightarrow{\underline{a}} \langle F', \sigma' \rangle}{\langle E \boxtimes_M F, \sigma \rangle \xrightarrow{\underline{a}} \langle E \boxtimes_M F', \sigma' \rangle} (\underline{a} \notin M)$
Cooperation with synchronisation:	$\frac{\langle E, \sigma \rangle \xrightarrow{\underline{a}} \langle E', \tau \rangle \quad \langle F, \sigma \rangle \xrightarrow{\underline{a}} \langle F', \tau' \rangle}{\langle E \boxtimes_M F, \sigma \rangle \xrightarrow{\underline{a}} \langle E' \boxtimes_M F', \Gamma(\sigma, \tau, \tau') \rangle} (\underline{a} \in M, \Gamma \text{ defined})$
Constant:	$\frac{\langle E, \sigma \rangle \xrightarrow{\underline{a}} \langle E', \sigma' \rangle}{\langle A, \sigma \rangle \xrightarrow{\underline{a}} \langle E', \sigma' \rangle} (A \stackrel{def}{=} E)$

Fig. 2. Operational semantics for HYPE

$$\frac{\langle Heat, \sigma \rangle \xrightarrow{\text{init}} \langle Heat, \sigma_1 \rangle \quad \langle Shade, \sigma \rangle \xrightarrow{\text{init}} \langle Shade, \sigma_3 \rangle}{\langle Heat \boxtimes_{\text{init}} Shade, \sigma \rangle \xrightarrow{\text{init}} \langle Heat \boxtimes_{\text{init}} Shade, \sigma_3 \rangle}$$

$$\begin{aligned} \sigma &= \{h \mapsto *, \quad d \mapsto *, \quad s \mapsto *, c \mapsto *, t \mapsto *\} \\ \sigma_1 &= \sigma[h \mapsto (0, const)] = \{h \mapsto (0, const), d \mapsto *, \quad s \mapsto *, c \mapsto *, t \mapsto *\} \\ \sigma_2 &= \sigma[d \mapsto (0, const)] = \{h \mapsto *, \quad d \mapsto (0, const), s \mapsto *, c \mapsto *, t \mapsto *\} \\ \sigma_3 &= \Gamma(\sigma, \sigma_1, \sigma_2) = \{h \mapsto (0, const), d \mapsto (0, const), s \mapsto *, c \mapsto *, t \mapsto *\} \end{aligned}$$

Fig. 3. Example derivation using Γ (* indicates an undefined value)

form as an activity to reflect the fact that the state captures the activities that are currently in effect. The notion of state here is not a valuation of system variables but rather a collection of flows that occur in the system.

The operational semantics give a labelled transition system over configurations $(\mathcal{F}, \mathcal{E}, \rightarrow \subseteq \mathcal{F} \times \mathcal{E} \times \mathcal{F})$. We write $F \xrightarrow{\underline{a}} F'$ for $(F, \underline{a}, F') \in \rightarrow$. In the following, $E, F \in \mathcal{C}_{Sys}$. The rules are given in Figure 2 and are fairly standard. In Choice, Prefix without influence, Cooperation without synchronisation and Constant, states are not changed by the application of the rule. For Prefix with influence, the state needs to be updated, and for Cooperation with synchronisation, the two new states in the premise of the rule are merged using

$$\begin{aligned}
\sigma_0 &= \{h \mapsto (0, \text{const}), d \mapsto (0, \text{const}), s \mapsto (0, \text{const}), c \mapsto (-1, \text{linear}(K)), t \mapsto (1, \text{const})\} \\
\sigma_1 &= \{h \mapsto (0, \text{const}), d \mapsto (0, \text{const}), s \mapsto (r_s, \text{const}), c \mapsto (-1, \text{linear}(K)), t \mapsto (1, \text{const})\} \\
\sigma_2 &= \{h \mapsto (0, \text{const}), d \mapsto (-r_d, \text{const}), s \mapsto (0, \text{const}), c \mapsto (-1, \text{linear}(K)), t \mapsto (1, \text{const})\} \\
\sigma_3 &= \{h \mapsto (0, \text{const}), d \mapsto (-r_d, \text{const}), s \mapsto (r_s, \text{const}), c \mapsto (-1, \text{linear}(K)), t \mapsto (1, \text{const})\} \\
\sigma_4 &= \{h \mapsto (r_h, \text{const}), d \mapsto (0, \text{const}), s \mapsto (0, \text{const}), c \mapsto (-1, \text{linear}(K)), t \mapsto (1, \text{const})\} \\
\sigma_5 &= \{h \mapsto (r_h, \text{const}), d \mapsto (0, \text{const}), s \mapsto (r_s, \text{const}), c \mapsto (-1, \text{linear}(K)), t \mapsto (1, \text{const})\} \\
\sigma_6 &= \{h \mapsto (r_h, \text{const}), d \mapsto (-r_d, \text{const}), s \mapsto (0, \text{const}), c \mapsto (-1, \text{linear}(K)), t \mapsto (1, \text{const})\} \\
\sigma_7 &= \{h \mapsto (r_h, \text{const}), d \mapsto (-r_d, \text{const}), s \mapsto (r_s, \text{const}), c \mapsto (-1, \text{linear}(K)), t \mapsto (1, \text{const})\}
\end{aligned}$$

Fig. 4. The states of the orbiter temperature control system

the function Γ . The updating function $\sigma[\iota \mapsto (r, I)]$ is defined as

$$\sigma[\iota \mapsto (r, I)](x) = \begin{cases} (r, I) & \text{if } x = \iota \\ \sigma(x) & \text{otherwise.} \end{cases}$$

The notation $\sigma[u]$ will also be used for an update, with $\sigma[u_1 \dots u_n]$ denoting $\sigma[u_1] \dots [u_n]$. The partial function $\Gamma : \mathcal{S} \times \mathcal{S} \times \mathcal{S} \rightarrow \mathcal{S}$ is defined as follows.

$$(\Gamma(\sigma, \tau, \tau'))(\iota) = \begin{cases} \tau(\iota) & \text{if } \sigma(\iota) = \tau'(\iota), \\ \tau'(\iota) & \text{if } \sigma(\iota) = \tau(\iota), \\ \text{undefined} & \text{otherwise.} \end{cases}$$

When synchronisation occurs, two states must be merged and the function uses the previous state and the new states to determine which values have changed and then puts these changed values into the new state. Γ will be undefined if both the second and third argument differ from the first argument, namely if the values in the new state both differ from the old state since this represents conflicting updates. Figure 3 illustrates the use of Γ . We show in Section 5 how we can constrain the syntactic form of our models to ensure that conflicts do not occur in updates while also ensuring well-structured models.

The next two definitions will be useful in referring to the states of the model.

Definition 3.2. The *derivative set* of a controlled system P , $\text{ds}(P)$ is defined as the smallest set satisfying

- if $\langle P, \sigma \rangle \xrightarrow{\text{init}} \langle P', \sigma' \rangle$ then $\langle P', \sigma' \rangle \in \text{ds}(P)$
- if $\langle P', \sigma' \rangle \in \text{ds}(P)$ and $\langle P', \sigma' \rangle \xrightarrow{a} \langle P'', \sigma'' \rangle$ then $\langle P'', \sigma'' \rangle \in \text{ds}(P)$.

Definition 3.3. The set of *states* of the derivative set of a controlled system P is defined as $\text{st}(P) = \{\sigma \mid \langle Q, \sigma \rangle \in \text{ds}(P)\}$.

In the *TempCtrl* model, there are eight states of interest (we have omitted the state before the *init* event) given in Figure 4. Each state captures the influences that are currently active. Since the influence strengths and types of c and t do not change, and each of h , d and s have two possible strengths, there are eight states. For example, σ_3 reflects that the heater (h) is off (and has no effect), the shade (d) is up, the sun (s) is shining, cooling (c) is happening, and time (t) is passing.

4. Hybrid Semantics

We have defined the labelled transition system via the operational semantics and it captures the discrete behaviour of the system in terms of events. We now move on to consider the continuous behaviour of the system. We extract a set of ODEs for each state which appears in a configuration in the labelled transition system, thus describing the continuous behaviour at that configuration. We will label this set as P_σ where P is the constant used for the controlled system and σ is the state.

Given a controlled system P , and a derivative $\langle P', \sigma \rangle \in \text{ds}(P)$, the ODEs associated with the state σ are

defined as follows.

$$P'_\sigma = \left\{ \frac{dV}{dt} = \sum \{ r \times \llbracket I(\mathcal{W}) \rrbracket \mid iv(\iota) = V \text{ and } \sigma(\iota) = (r, I(\mathcal{W})) \} \mid V \in \mathcal{V} \right\}$$

For each variable, we identify the influences that are associated with that variable, and then, using the information in the state, the ODE is constructed from the product of the strength and type of each influence. The influence type allows the variables on which the variable depend to appear in the ODE. Hence, for each state, we have a collection of ODEs, one for each variable V . We keep influence strength and influence type separate since in many examples, the influence strength varies whereas the influence type does not, allowing a fine-grained description of dynamics. For a discussion of the range of ODEs expressible in HYPE, refer to Section 9.

We obtain the following ODE from σ_3 .

$$TempCtrl_{\sigma_3} = \left\{ \frac{dK}{dt} = r_s - r_d - K, \frac{dT}{dt} = 1 \right\}$$

Therefore, this process enables us to obtain ODEs describing how the continuous part of the system evolves, and we have different sets of ODEs to describe the different dynamics that can be in operation. We wish to combine this information with the event conditions already defined and an obvious way to do this is to translate this information into a hybrid automaton. Therefore, this well-supported formalism provides a powerful back-end for HYPE.

Hybrid automata are dynamic systems presenting both discrete and continuous evolution. They consist of variables evolving continuously in time, subject to abrupt changes induced by discrete instantaneous *control* events. When discrete events happen the automaton enters its next *mode*, where the rules governing the flow of continuous variables change. See [Hen96] for further details.

Definition 4.1. A *hybrid automaton* is a tuple $\mathcal{H} = (V, E, \mathbf{X}, \mathcal{E}, flow, init, inv, event, jump, reset, urgent)$, where

- $\mathbf{X} = \{X_1, \dots, X_n\}$ is a finite set of real-valued variables. The time derivative of X_j is \dot{X}_j , and the value of X_j after a change of *mode* is X'_j .
- The *control graph* $G = (V, E)$ is a finite labelled multigraph. Vertices in the set V are the (*control*) *modes*, while edges in the multiset E are called (*control*) *switches* and model the happening of a discrete event.
- Associated with each vertex $v \in V$ there is a set of ordinary differential equations $\dot{\mathbf{X}} = flow(v)$ referred to as the *flow conditions*, where $flow(v)$ is a vector field from \mathbb{R}^n to \mathbb{R}^n . Moreover, $init(v)$ and $inv(v)$ are two formulae on \mathbf{X} specifying the *admissible initial conditions* and some *invariant conditions* that must be true during the continuous evolution of variables in v .
- Edges $e \in E$ of the control graph are labelled by an event $event(e) \in \mathcal{E}$ and by $jump(e)$, a formula on \mathbf{X} stating for which values of variables each transition is active, and by $reset(e)$, a formula on $\mathbf{X} \cup \mathbf{X}'$ specifying the change of the variables' values after the transition has taken place. Moreover, each edge $e \in E$ can be declared urgent, by setting to true the boolean flag $urgent(e)$, meaning that the transition is taken at once when $jump(e)$ becomes true. Otherwise, the transition can be taken nondeterministically whenever $jump(e)$ is true. Two edges with the same vertices must differ in at least one of *event*, *jump* or *reset*.

Definition 4.2. Given a hybrid automaton \mathcal{H} , an (*automaton*) *state* is a pair (v, \mathbf{x}) for $v \in V$ and $\mathbf{x} = (x_1, \dots, x_n)$ a tuple of values in \mathbb{R}^n .

Consider a HYPE model $(P_0, \mathcal{V}, IN, IT, \mathcal{E}, \mathcal{A}, ec, iv, EC, ID)$ with initial configuration $\langle P_0, \sigma_0 \rangle \in \mathcal{F}$. For P_0 the only possible transition is the event init. Let $\langle P, \sigma \rangle$ be the configuration reached after its occurrence, $\langle P_0, \sigma_0 \rangle \xrightarrow{\text{init}} \langle P, \sigma \rangle$. We introduce some notation for ease of referring to event conditions.

Definition 4.3. Given a HYPE model with $ec : \mathcal{E} \rightarrow EC$, let $ec(\underline{a}) = (act(\underline{a}), res(\underline{a}))$ where $act(\underline{a})$ describes the activation condition for event \underline{a} and $res(\underline{a})$ which describes the reset for event \underline{a} .

Definition 4.4. The hybrid automaton $\mathcal{H} = (V, E, \mathbf{X}, \mathcal{E}, flow, init, inv, event, jump, reset, urgent)$ can be obtained from the HYPE model $(P_0, \mathcal{V}, IN, IT, \mathcal{E}, \mathcal{A}, ec, iv, EC, ID)$ as follows.

- The set of modes V is the set of configurations reachable in 0 or more steps from $\langle P, \sigma \rangle$, namely $ds(P_0)$.

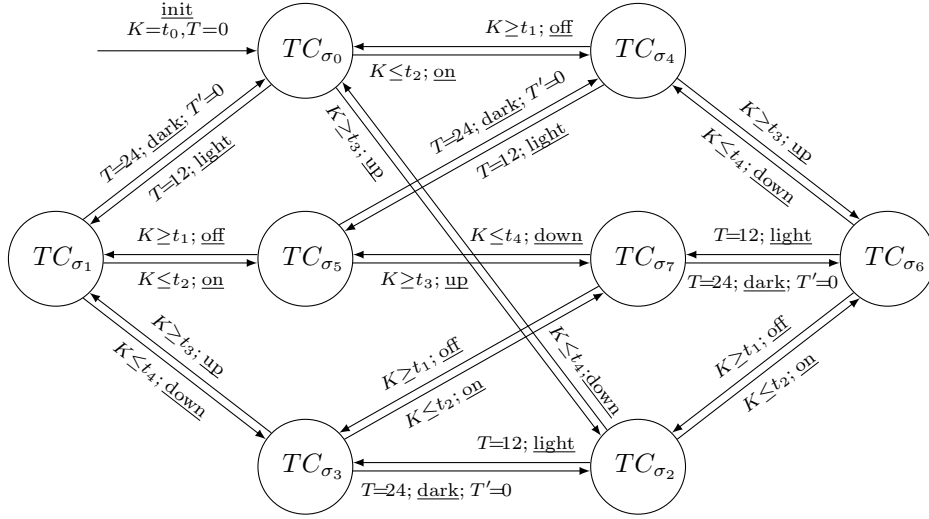


Fig. 5. Hybrid automaton of the orbiter temperature control system

- The edges E of the control graph connect two modes (v_1, v_2) iff $v_1 = \langle P_1, \sigma_1 \rangle$, $v_2 = \langle P_2, \sigma_2 \rangle$ and $\langle P_1, \sigma_1 \rangle \xrightarrow{\underline{a}} \langle P_2, \sigma_2 \rangle$ is a derivation for some \underline{a} .
- $\mathbf{X} = \mathcal{V}$ is the set of variables of the HYPE system.
- \mathcal{E} is the set of events \mathcal{E} of P_0 .
- Let $v_j = \langle P_j, \sigma_j \rangle$, then $flow(v_j)[X_i] = \sum \{r[I(\mathcal{W})] \mid iv(\iota) = X_i \text{ and } \sigma_j(\iota) = (r, I(\mathcal{W}))\}$.
- $init(v) = \begin{cases} res(\underline{init}), & \text{if } v = \langle P, \sigma \rangle \\ false, & \text{otherwise} \end{cases}$ with primes removed from variables⁵.
- $inv(v) = true$.
- Let $e = (\langle P_1, \sigma_1 \rangle, \langle P_2, \sigma_2 \rangle)$ with $\langle P_1, \sigma_1 \rangle \xrightarrow{\underline{a}} \langle P_2, \sigma_2 \rangle$. Then $event(e) = \underline{a}$ and $reset(e) = res(\underline{a})$. Moreover, if $act(\underline{a}) \neq \perp$, then $jump(e) = act(\underline{a})$ and $urgent(e) = true$, otherwise $jump(e) = true$ and $urgent(e) = false$.

Given a HYPE model P , we denote its hybrid automaton as $\mathcal{H}(P)$. As long as there is a finite number of subcomponents (since infinite sum is not possible) and a finite number of controller components, then the number of modes is finite⁶. Figure 5 shows the HYPE model from the example as a hybrid automaton. Note that in this case, which modes are visited is determined by the values of k_i , $i = 0, \dots, 4$.

We do not use invariants in HYPE models because they are specific to modes of the hybrid automata, which are obtained in HYPE as the possible configurations of the controlled system. This implies that invariants should be specified within the controlled system, not orthogonally as for event conditions. This conflicts with the modelling philosophy of HYPE. However, the behaviour of invariants can be obtained in our setting by incorporating them in the activation condition of events, that is by forcing events to happen as soon as the conditions that would have been specified by an invariant become false.

The behaviour of a hybrid automaton can be defined in terms of a transition system [Hen96].

Definition 4.5. The *timed transition system* of a hybrid automaton has a set of states $\mathcal{S} \subseteq V \times \mathbb{R}^n$ with initial states \mathcal{S}_0 and transitions labelled with elements from $\mathcal{E} \cup \mathbb{R}$ such that

1. $(v, \mathbf{x}) \in \mathcal{S}$ whenever $inv(v)(\mathbf{x})$ holds and $(v, \mathbf{x}) \in \mathcal{S}_0$ whenever $init(v)(\mathbf{x})$ holds.
2. $(v, \mathbf{x}) \xrightarrow{a} (v', \mathbf{x}')$ whenever there is an edge $e \in E$ such that $e = (v, v')$, $event(e) = a$, $jump(e)(\mathbf{x})$ and $reset(e)(\mathbf{x}, \mathbf{x}')$ both hold.

⁵ $res(\underline{init})$ is a reset so uses primed variables to refer to the new values of variables whereas $init(v)$ is an initialisation condition and refers to variables without primes.

⁶ To ensure finiteness of controller components, recurrent controller components with increasing indices should be avoided, for example $Con_i \stackrel{def}{=} \underline{a}.Con_{i+1}$. This type of unbounded controller is unlikely to be required when modelling hybrid systems, and hence a maximum controller can be specified, for example $Con_n \stackrel{def}{=} \underline{a}.0$.

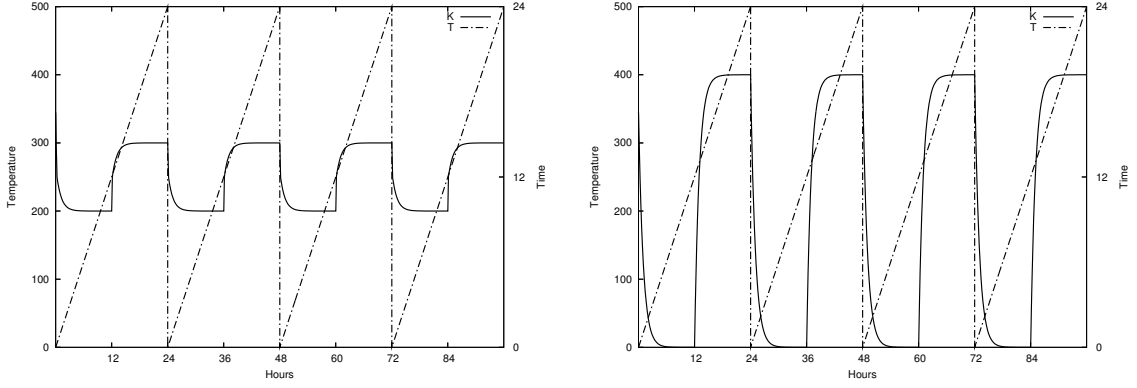


Fig. 6. Traces for the orbiter with shade and heating (left) and without shade and heating (right) where $r_h = 200$, $r_d = 100$, $r_s = 400$, $k_0 = 350$, $k_1 = 250$, $k_2 = 250$, $k_3 = 300$ and $k_4 = 300$.

3. $(v, \mathbf{x}) \xrightarrow{\tau} (v', \mathbf{x}')$ whenever $v = v'$ and there is a differentiable function $f : [0, \tau] \rightarrow \mathbb{R}^n$ with first derivative $f' : (0, \tau) \rightarrow \mathbb{R}^n$ such that $f(0) = \mathbf{x}$ and $f(\tau) = \mathbf{x}'$ and for all $t \in (0, \tau)$, $inv(v)(\mathbf{x})$ holds and $flow(v)(\mathbf{x}) = f'(\mathbf{x})$. Moreover, for all $t \in (0, \tau)$ and all $e = (v, u) \in E$ with $urgent(e) = true$, $jump(e)(f(t)) = false$.

The third condition of this definition ensures that a transition associated with the passing of time can only proceed if no jump condition becomes true for an urgent edge. Therefore continuous behaviour can only continue if a mode change is not required. Note that for our mapping of HYPE models to hybrid automata, this definition is somewhat simplified by the lack of invariants.

Definition 4.6. Given a timed transition of a hybrid automaton, a *trace* of the hybrid automaton is a sequence of the form $(v_0, \mathbf{x}_0) \xrightarrow{\rho_0} (v_1, \mathbf{x}_1) \xrightarrow{\rho_1} \dots$ where $(v_i, \mathbf{x}_i) \in \mathcal{S}$, $(v_0, \mathbf{x}_0) \in \mathcal{S}_0$ and $\rho_i \in \mathcal{E} \cup \mathbb{R}$.

With the appropriate software, we can construct simulations of our HYPE models, and demonstrate their traces graphically. Note that HYPE models are not necessarily deterministic. The presence of non-urgent transitions can introduce non-determinism, as can choice between events when activation conditions become true in the same time instant. In the case of the orbiter, we do have determinism. A trace is demonstrated in the left graph of Figure 6 for the parameters $r_h = 200$, $r_d = 100$, $r_s = 400$, $k_1 = 250$, $k_2 = 250$, $k_3 = 300$, $k_4 = 300$ and $k_0 = 350$. If we remove the temperature controls in the orbiter, we obtain the right hand graph which demonstrates the temperature extremes which can be achieved. This graph is based on the HYPE model

$$((Sun \boxtimes_{\{init\}} Cool(K)) \boxtimes_{\{init, light, dark\}} Time) \boxtimes_{\{init, light, dark\}} init.Con_s$$

This section has defined the semantics of HYPE models, both discrete and continuous. In the example, the subcomponents of the orbiter have a very specific form which is not required by the HYPE definition which is more liberal. In the next two section, we consider how constraining some aspects of models can result in models with desirable behaviour.

5. Well-defined HYPE models

Now that we have defined the behaviour of HYPE models, we wish to specify a subset of HYPE models that have “good” structure and behaviour. We introduce the notion of a well-defined HYPE model and reason about how this notion affects derivation of transitions specifically where the cooperation operator is used. First we require some auxiliary definitions.

Definition 5.1. Given a controlled system P , the *set of events*, $ev(P)$ is defined structurally as

$$ev(\underline{a} : \alpha.S) = \{\underline{a}\} \quad ev(\underline{a}.S) = \{\underline{a}\} \quad ev(S_1 + S_2) = ev(S_1) \cup ev(S_2) \quad ev(P_1 \boxtimes_L P_2) = ev(P_1) \cup ev(P_2)$$

Definition 5.2. Given a controlled system P , the *set of influences*, $in(P)$ is defined structurally as

$$in(\underline{a} : (\iota, r, I(\mathcal{W})).S) = \{\iota\} \quad in(\underline{a}.S) = \emptyset \quad in(S_1 + S_2) = in(S_1) \cup in(S_2) \quad in(P_1 \boxtimes_L P_2) = in(P_1) \cup in(P_2)$$

Definition 5.3. Given a controlled system P , the *set of prefixes*, $\text{pr}(P)$ is defined structurally as

$$\text{pr}(\underline{a}:\alpha.S) = \{\underline{a}:\alpha\} \quad \text{pr}(\underline{a}.S) = \emptyset \quad \text{pr}(S_1 + S_2) = \text{pr}(S_1) \cup \text{pr}(S_2) \quad \text{pr}(P_1 \underset{L}{\bowtie} P_2) = \text{pr}(P_1) \cup \text{pr}(P_2)$$

Definition 5.4.

A *well-defined* controlled system has the following properties.

1. Subcomponents have the form $S(\mathcal{W}) \stackrel{\text{def}}{=} \sum_{j=1}^n \underline{a}_j:(\iota, r_j, I_j(\mathcal{W})).S(\mathcal{W}) + \underline{\text{init}}:(\iota, r, I(\mathcal{W})).S(\mathcal{W})$ where $n \geq 0$, $\underline{a}_j \neq \underline{a}_k$ for $j \neq k$ and $\underline{a}_j \neq \underline{\text{init}}$ for all j .
2. For each pair of subcomponents S_i and S_j , $\text{in}(S_i) \cap \text{in}(S_j) = \emptyset$ for $i \neq j$.
3. For a component $C_i \underset{L}{\bowtie} C_j$, $\text{ev}(C_i) \cap \text{ev}(C_j) = L$.
4. For the uncontrolled system Σ , $\text{ev}(\Sigma) = \mathcal{E}$ and $\text{in}(\Sigma) = IN$.
5. For the controlled system $\Sigma \underset{L}{\bowtie} \underline{\text{init}}.Con$, $\text{ev}(\Sigma) \cap \text{ev}(\underline{\text{init}}.Con) = L$ and $\text{ev}(Con) \subseteq \text{ev}(\Sigma)$.

This definition ensures that HYPE models are built up out of individual flows in a consistent manner. Each subcomponent captures the events that affect a particular influence and this influence is linked to a variable by the function iv . The first two conditions ensure a one-to-one relationship between subcomponents and influence names. The third and fifth conditions ensure that when an event happens, all subcomponents and controllers react to it. We use the symbol $\underset{*}{\bowtie}$ in HYPE models to show when these two conditions hold in a model. The fourth condition ensures there are no unused events or influence names, and is useful to know when defining the synchronisation of two HYPE models. The orbiter is a well-defined HYPE model.

The fifth condition permits events to appear in the uncontrolled system that do not appear in the controller, since not all events need to depend causally on other events, and we do not wish to compel modellers to define a controller for each of these events. However, for simplicity of definitions, in the remainder of this document, we will assume that there is a controller for each event \underline{u} that does not appear in any defined controller, with the form $Con_{\underline{u}} \stackrel{\text{def}}{=} \underline{u}.Con_{\underline{u}}$.

Not only does well-definedness ensure our models have a specific form that exemplifies how we build them, but as we will show in Section 8 this form is useful when we consider equivalence semantics and reason about the behaviour of our models. Note that we also impose, through the definition of a HYPE model, the conditions that resets and initial values are deterministic, namely a single value is defined. Additionally, activation conditions define closed sets, as guaranteed by the use of nonstrict inequalities and of positive boolean formulae to combine them.

We now present a result that follows from well-definedness. The function Γ used in the cooperation rules, is always defined. This ensures that it is always possible to derive a transition for two components in cooperation.

Proposition 5.1. Let E and F be components of a well-defined HYPE model such that $\langle E, \sigma \rangle \xrightarrow{\underline{a}} \langle E', \tau \rangle$ and $\langle F, \sigma \rangle \xrightarrow{\underline{a}} \langle F', \tau' \rangle$ then the transition $\langle E \underset{M}{\bowtie} F, \sigma \rangle \xrightarrow{\underline{a}} \langle E' \underset{M}{\bowtie} F', \Gamma(\sigma, \tau, \tau') \rangle$ is always possible for $\underline{a} \in M$.

Proof. The first two conditions of well-definedness ensure that each subcomponent is associated with exactly one influence name and *vice versa*. Therefore $\text{in}(E) \cap \text{in}(F) = \emptyset$. Let ι be an influence name appearing in E . Since it does not appear in F , $\sigma(\iota) = \tau'(\iota)$ since it cannot have been modified in the derivation of F , hence $(\Gamma(\sigma, \tau, \tau'))(\iota)$ is defined. A similar argument can be applied to influence names appearing in F and not E . Hence $\Gamma(\sigma, \tau, \tau')$ is defined. \square

Additionally, we provide a result about the form that the configurations of a HYPE model take.

Proposition 5.2. If P is a well-defined HYPE model with controlled system $P \stackrel{\text{def}}{=} \Sigma \underset{*}{\bowtie} \underline{\text{init}}.Con$ then $ds(P) = \{\langle \Sigma \underset{*}{\bowtie} Con, \tau \rangle, \langle \Sigma \underset{*}{\bowtie} D_1, \sigma_1 \rangle, \dots, \langle \Sigma \underset{*}{\bowtie} D_n, \sigma_n \rangle\}$ where, for arbitrary state σ , the derivative set of the controller is $ds(\underline{\text{init}}.Con) = \{\langle Con, \sigma \rangle, \langle D_1, \sigma \rangle, \dots, \langle D_n, \sigma \rangle\}$. Moreover, each $\Sigma \underset{*}{\bowtie} D_i$ is well-defined.

Proof. First note that by the particular form of Σ imposed by well-definedness, for any event \underline{a} that appears in Σ , $\langle \Sigma, \rho \rangle \xrightarrow{\underline{a}} \langle \Sigma, \rho' \rangle$ for some states ρ, ρ' . This means that the configurations of $\Sigma \underset{*}{\bowtie} Con$ are totally determined by the events that are allowed by the controller. Moreover, the value of σ in $ds(Con)$ does not determine these events. Additionally, $\underset{*}{\bowtie}$ is a static operator. For well-definedness, note that Conditions 1-4 in Definition 5 hold for Σ since $\Sigma \underset{*}{\bowtie} \underline{\text{init}}.Con$ is well-defined, and Condition 5 is only required to hold of the system before the event $\underline{\text{init}}$ has occurred. \square

We can also reason about derivatives and the relationship between states. In well-defined HYPE models, subcomponents are defined as a number of simple loops, reflecting events that can occur and the associated changes in the continuous part of the system. For such systems, we can show that the current state in a configuration does not determine which future events happen (only the controller can influence this) and hence the state can be discounted in certain settings.

The next proposition shows how, if we know how one state is related to a prior state, then we can conclude that a similar relationship holds if we use a different state. Additionally, we can express a state modification in terms of the a minimum number of updates, one for each influence that changes.

Proposition 5.3. In a well-defined controlled system, if $\langle P', \sigma[l_1 \mapsto (r_1, I_1)] \dots [l_n \mapsto (r_n, I_n)] \rangle$ is a derivative of $\langle P, \sigma \rangle$, then $\langle P', \tau[l_1 \mapsto (r_1, I_1)] \dots [l_n \mapsto (r_n, I_n)] \rangle$ is a derivative of $\langle P, \tau \rangle$. Moreover, it is the case that $\sigma[l_1 \mapsto (r_1, I_1)] \dots [l_n \mapsto (r_n, I_n)] = \sigma[l'_1 \mapsto (r'_1, I'_1)] \dots [l'_m \mapsto (r'_m, I'_m)]$ for $m \leq n$ where each influence name l'_k appears at most once.

Proof. We first consider the single step case and prove this by induction on the the height of the derivation of the transition. Symmetrical cases have been omitted.

Prefix with influence $P \equiv \underline{a} : (\iota, r, I).P_1$ so $\langle \underline{a} : (\iota, r, I).P_1, \sigma \rangle \xrightarrow{\underline{a}} \langle P_1, \sigma[l \mapsto (r, I)] \rangle$. Similarly, we also have $\langle \underline{a} : (\iota, r, I).P_1, \tau \rangle \xrightarrow{\underline{a}} \langle P_1, \tau[l \mapsto (r, I)] \rangle$.

Prefix without influence $P \equiv \underline{a}.P_1$ so $\langle \underline{a}.P_1, \sigma \rangle \xrightarrow{\underline{a}} \langle P_1, \sigma \rangle$. Similarly, $\langle \underline{a}.P_1, \tau \rangle \xrightarrow{\underline{a}} \langle P_1, \tau \rangle$.

Choice $P \equiv P_1 + P_2$ so $\langle P_1, \sigma \rangle \xrightarrow{\underline{a}} \langle P', \sigma[l_1 \mapsto (r_1, I_1)] \dots [l_n \mapsto (r_n, I_n)] \rangle$ by a shorter inference. By the inductive hypothesis, it is the case that $\langle P_1, \tau \rangle \xrightarrow{\underline{a}} \langle P', \tau[l_1 \mapsto (r_1, I_1)] \dots [l_n \mapsto (r_n, I_n)] \rangle$, and therefore $\langle P_1 + P_2, \tau \rangle \xrightarrow{\underline{a}} \langle P', \tau[l_1 \mapsto (r_1, I_1)] \dots [l_n \mapsto (r_n, I_n)] \rangle$.

Cooperation without synchronisation $P \equiv P_1 \underset{L}{\bowtie} P_2$ with $\underline{a} \notin L$ so we know that $\langle P_1, \sigma \rangle \xrightarrow{\underline{a}} \langle P', \sigma[l_1 \mapsto (r_1, I_1)] \dots [l_n \mapsto (r_n, I_n)] \rangle$ by a shorter inference. Straightforward.

Cooperation with synchronisation $P \equiv P_1 \underset{L}{\bowtie} P_2$ with $\underline{a} \in L$. We have the transitions $\langle P_1, \sigma \rangle \xrightarrow{\underline{a}} \langle P'_1, \sigma'_1 \rangle$ and $\langle P_2, \sigma \rangle \xrightarrow{\underline{a}} \langle P'_2, \sigma'_2 \rangle$ by a shorter inference where $\sigma'_1 = \sigma[\eta_1 \mapsto (s_1, J_1)] \dots [\eta_m \mapsto (s_m, J_m)]$ and $\sigma'_2 = \sigma[\xi_1 \mapsto (t_1, K_1)] \dots [\xi_p \mapsto (t_p, K_p)]$ such that $\Gamma(\sigma, \sigma'_1, \sigma'_2) = \sigma[l_1 \mapsto (r_1, I_1)] \dots [l_n \mapsto (r_n, I_n)]$ with $P' \equiv \langle P'_1 \underset{L}{\bowtie} P'_2, \Gamma(\sigma, \sigma'_1, \sigma'_2) \rangle$. So by the inductive hypothesis $\langle P_1, \tau \rangle \xrightarrow{\underline{a}} \langle P'_1, \tau'_1 \rangle$ and $\langle P_2, \tau \rangle \xrightarrow{\underline{a}} \langle P'_2, \tau'_2 \rangle$ where $\tau'_1 = \tau[\eta_1 \mapsto (s_1, J_1)] \dots [\eta_m \mapsto (s_m, J_m)]$ and $\tau'_2 = \tau[\xi_1 \mapsto (t_1, K_1)] \dots [\xi_p \mapsto (t_p, K_p)]$ and $\langle P, \tau \rangle \xrightarrow{\underline{a}} \langle P'_1 \underset{L}{\bowtie} P'_2, \Gamma(\tau, \tau'_1, \tau'_2) \rangle$. Note that $\Gamma(\tau, \tau'_1, \tau'_2) = \tau[l_1 \mapsto (r_1, I_1)] \dots [l_n \mapsto (r_n, I_n)]$ since the same influences have been updated to the same values, and Γ is defined piecewise over influences.

Constant Straightforward.

We can use an identical proof to go from an $(n - 1)$ -step derivation to an n -step derivation. Hence by induction, over the length of the derivation, this result holds for all derivations.

Finally, because of the lack of dependence on previous states in the labelled transition system, it is possible to reduce the number of updates by only retaining the last update for a given influence name ι . \square

The next proposition shows that the state is irrelevant when considering the actions that can be performed by a system.

Proposition 5.4. The labelled transition systems of $\langle P, \sigma \rangle$ and $\langle P, \tau \rangle$ are isomorphic if P is well-defined.

Proof. This can be proved using Proposition 5.3. \square

We can also consider what happens to a system after the first init event. Note that the definition of a well-defined controlled system requires that the first transition be an init event since all events must appear in the cooperation set and *Con* is prefixed by init. The following proof shows that the starting state is irrelevant since the init action will set every value in the state.

Proposition 5.5. Let P be a well-defined controlled system. If $\langle P, \sigma \rangle \xrightarrow{\text{init}} \langle P', \sigma' \rangle$ and $\langle P, \tau \rangle \xrightarrow{\text{init}} \langle P', \tau' \rangle$ then $\sigma' = \tau'$.

Proof. By definition of a well-defined controlled system, there is a prefix containing init and ι for every ι in the controlled system. Assuming $\mathcal{A} = \{\iota_1, \dots, \iota_n\}$ which appear in the controlled system and no others, then a state σ maps each ι to a rate and influence type pair. After an init each ι will be updated and the state before the transition is irrelevant, namely $\sigma' = \sigma[l_1 \mapsto (r_1, I_1)] \dots [l_n \mapsto (r_n, I_n)] = \emptyset[l_1 \mapsto (r_1, I_1)] \dots [l_n \mapsto (r_n, I_n)] = \tau[l_1 \mapsto (r_1, I_1)] \dots [l_n \mapsto (r_n, I_n)] = \tau'$ by Proposition 5.3. \square

6. Well-behaved HYPE models

We have considered the syntactic form that HYPE models can take, and the advantages of this at the level of the labelled transition system. We now consider a notion of behaviour that ensures our models show reasonable behaviour, and motivate it by considering the underlying hybrid automaton of the model.

Definition 6.1. A HYPE model P is well-behaved if it has a finite number of finite sequences of simultaneous instantaneous events and these sequences are independent of the initial state of the system.

This ensures that our models cannot perform an infinite number of discrete steps in a single time instant. We call this unwanted behaviour *instantaneous Zeno behaviour* by analogy with Zeno behaviour which describes the possibility of an infinite number of steps in a finite time period. Our focus on this unwanted instantaneous behaviour is informed by our later work [BGH10a] where we consider Piecewise Deterministic Markov Processes [Dav93] as a semantic model which excludes this behaviour.

We now characterise well-behaved HYPE models in terms of their event conditions and controllers by constructing an instantaneous activation graph or I-graph for a model, and checking it for acyclicity.

The I-graph captures information about the flow of data in the HYPE model. It simulates the behaviour of the instantaneous transitions so that each sequence of simultaneous instantaneous events corresponds to a path in the graph. This overapproximates behaviour since it is based on constraints imposed by the event conditions and the controller only, and does not take account of changes that occur due to initial conditions or continuous flow. This means additional information that could further constrain the possible values is ignored. This is an appropriate approach to take because we want to reason about the HYPE model without investigating its full behaviour.

For each event, we consider the intersection between its activation region (the set of points making the activation condition true for the event) and its reset region (the set of points reachable immediately after the occurrence of the event). Such set-based constraints thus take into account overlap only at the set level and not at the single value level. This means that the I-graph does not capture the changes in the value of a variable, only the regions that these values can fall into. When necessary, we also overapproximate activation and reset regions with computable sets to ensure that testing for acyclicity of an I-graph is decidable.

We will show that if no cycles are found in the I-graph, then instantaneous Zeno behaviour is not possible in the HYPE model (which, in this context, means in the hybrid automaton), but we cannot know the converse – a cycle in the graph does not guarantee that there is instantaneous Zeno behaviour in the model since these are overapproximations.

We start with some preliminary definitions.

Definition 6.2. Given an event \underline{a} , $\mathcal{V}(\underline{a})$ is the set of variables that appear explicitly in $ec(\underline{a}) = (act(\underline{a}), res(\underline{a}))$.

Hence $\mathcal{V}(\underline{a})$ contains all of the variables in \mathcal{V} except those that have implicit identity resets for \underline{a} . From this, we can define when two events have disjoint (explicit) variables.

Definition 6.3. Events \underline{a} and \underline{b} are *independent* if $\mathcal{V}(\underline{a}, \underline{b}) = \mathcal{V}(\underline{a}) \cap \mathcal{V}(\underline{b}) = \emptyset$.

We wish to investigate the overlap of values between resets of one event and activation conditions of another, and also take into account the scheduling of events by the controller. Let $r_{\underline{a}}$ be the reset $res(\underline{a})$, expressed as a function.

Definition 6.4. For any event \underline{a} , let $G_{\underline{a}} = \{\mathbf{x} \in \mathbb{R}^n \mid act(\underline{a}) \text{ is true for } \mathbf{x}\}$ and $R_{\underline{a}} = r_{\underline{a}}(G_{\underline{a}})$

$G_{\underline{a}}$ describes all points at which $act(\underline{a})$ is true and hence the event \underline{a} is active, and $R_{\underline{a}}$ takes an image of this set, so we can consider $R_{\underline{a}}$ as all possible values that the system can take after the occurrence of \underline{a} . For $V \notin \mathcal{V}(\underline{a})$, the value associated with V will be \mathbb{R} . We can then determine enablers and inhibitors of events. In the following definition, $\pi_{\{V_{i_1}, \dots, V_{i_p}\}}(A) \subseteq \mathbb{R}^p$ for $A \subseteq \mathbb{R}^n$ is the projection of A to the coordinates associated with variables V_{i_1}, \dots, V_{i_p} .

Definition 6.5. Consider two non-independent events \underline{a} and \underline{b} of an HYPE model. Then

- \underline{a} is a (*potential*) *enabler* of \underline{b} if and only if $\pi_{\mathcal{V}(\underline{a}, \underline{b})}(R_{\underline{a}}) \cap \pi_{\mathcal{V}(\underline{a}, \underline{b})}(G_{\underline{b}}) \neq \emptyset$;
- \underline{a} is an *inhibitor* of \underline{b} if and only if $\pi_{\mathcal{V}(\underline{a}, \underline{b})}(R_{\underline{a}}) \cap \pi_{\mathcal{V}(\underline{a}, \underline{b})}(G_{\underline{b}}) = \emptyset$.




Fig. 7. Two tanks with separate inputs and joint output

This definition considers how the resets of one event can affect whether another event can happen immediately after it. We can use these ideas as the basis for constructing a graph that captures the possible sequences of instantaneous events.

As mentioned in the previous section, we assume a controller for every uncontrolled event \underline{u} of the form $Con_{\underline{u}} \stackrel{\text{def}}{=} \underline{u}.Con_{\underline{u}}$. It is not sufficient only to require that uncontrolled events inhibit themselves (an obvious condition to avoid infinite instantaneous sequences of the same event) since their event conditions may interact with the event conditions of other events, and hence we must consider these controllers from the beginning of the construction of the I-graph.

First, we let $ds(\underline{\text{init}}.Con) = \{C_1, \dots, C_m\}$ be the controller states (by applying the definition of ds in the obvious manner to the controller, and then ignoring the state which is identical for each controller derivative). The state $\underline{\text{init}}.Con$ is not in this set. Let the instantaneous events appearing in the controller be $\{\underline{a}_1, \dots, \underline{a}_k\}$ and $ev(C_i)$ is the subset of events that can be executed in mode C_i of the controller. Note that we are not working compositionally here. For a vector κ of k values, we write $\kappa[i \leftarrow u]$ to represent the i th position being updated with the value u . This vector will record which events are enabled and which are inhibited.

Definition 6.6. Given a HYPE model P with controller Con , the *instantaneous activation graph*, or *I-graph*, is given by $G_P(Con) = (V_P(Con), E_P(Con))$ where vertices are defined by

$$V_P(Con) = \bigcup_{C_i \in ds(\underline{\text{init}}.Con)} \left(\{C_i\} \times ev(C_i) \times \{0, 1\}^k \right) \subseteq \{C_1, \dots, C_k\} \times \{\underline{a}_1, \dots, \underline{a}_k\} \times \{0, 1\}^k.$$

The edge $(C_i, \underline{a}_i, \kappa_i) \rightarrow (C_j, \underline{a}_j, \kappa_j)$ belongs to $E_P(Con)$ if and only if

1. $\langle C_i, \sigma \rangle \xrightarrow{\underline{a}_i} \langle C_j, \sigma \rangle$ is a transition in the labelled transition system of the controller,
2. $\kappa_i[i] = 1$,
3. $\kappa_j[j] = 1$,
4. $\kappa_j = \kappa_i[k_1 \leftarrow 0] \dots [k_s \leftarrow 0][l_1 \leftarrow 1] \dots [l_r \leftarrow 1]$ where $\underline{a}_{k_1}, \dots, \underline{a}_{k_s}$ are the events for which \underline{a}_i is an inhibitor and $\underline{a}_{l_1}, \dots, \underline{a}_{l_r}$ are the events for which \underline{a}_i is an enabler.

The boolean vector κ records if an event is enabled or not. Hence, the second condition ensures that \underline{a}_i can fire and the third condition ensures that an edge is only added when it is possible for the next action to fire – this reduces the number of edges in the graph. The fourth condition encodes the updating of what is enabled. A node $(C, \underline{a}, \kappa)$ describes being in the state C with \underline{a} as the next action together with a vector describing which events are possible before the event \underline{a} has occurred. The I-graph represents information about the structure of the controller and the event conditions of a HYPE model.

We now introduce an example to illustrate how this construction works. Figure 7 shows two tanks that each have an input, and which are drained to a common pipe with a switch that allows liquid to be drained from either tank A or tank B . Once A is empty, the tap is switched for B to drain and *vice versa*. Each tank stops filling once it is full with capacity c_A or c_B , and a very small amount of liquid is lost reducing the level by a small amount, say l which is much smaller than either capacity. The initial state of the model is when tank A is draining and B is filling and the initial values are between empty and full. The HYPE model is now defined. Here we use \parallel for $\begin{array}{c} \times \\ \emptyset \end{array}$.

$$\begin{aligned}
In_A &\stackrel{def}{=} \underline{full}_A:(in_A, 0, c).In_A + \underline{empty}_B:(in_A, r_{in}, c).In_A + \underline{init}:(in_A, r_{in}, c).In_A \\
Out_A &\stackrel{def}{=} \underline{empty}_A:(out_A, 0, c).Out_A + \underline{empty}_B:(out_A, -r_{out}, c).Out_A + \underline{init}:(out_A, -r_{out}, c).Out_A \\
\\
In_B &\stackrel{def}{=} \underline{full}_B:(in_B, 0, c).In_B + \underline{empty}_A:(in_B, s_{in}, c).In_B + \underline{init}:(in_B, s_{in}, c).In_B \\
Out_B &\stackrel{def}{=} \underline{empty}_B:(out_B, 0, c).Out_B + \underline{empty}_A:(out_B, -s_{out}, c).Out_B + \underline{init}:(out_B, 0, c).Out_B \\
\\
C_A &\stackrel{def}{=} \underline{full}_A.C_A & C_B &\stackrel{def}{=} \underline{full}_B.C_B \\
C_{dA} &\stackrel{def}{=} \underline{empty}_A.C_{dA} & C_{dB} &\stackrel{def}{=} \underline{empty}_B.C_{dB} \\
\\
System &\stackrel{def}{=} ((In_A \bowtie_N Out_A) \bowtie_{N \cup M} (In_B \bowtie_M Out_B)) & N &= \{\underline{empty}_B, \underline{init}\} \\
&\stackrel{def}{=} \bowtie_L \underline{init}.(C_A \parallel C_B \parallel C_{dA}) & M &= \{\underline{empty}_A, \underline{init}\} \\
& & L &= \{\underline{full}_A, \underline{full}_B\} \cup N \cup M \\
\\
iw(\iota_A) &= A & iw(\iota_B) &= B \\
\\
ec(\underline{empty}_A) &= (A = 0, true) & ec(\underline{full}_A) &= (A = c_A, A' = c_A - l) \\
ec(\underline{empty}_B) &= (B = 0, true) & ec(\underline{full}_B) &= (B = c_B, B' = c_B - l) \\
ec(\underline{init}) &= (true, A' = A_0 \wedge B' = B_0) & \llbracket c \rrbracket &= 1
\end{aligned}$$

Fig. 8. HYPE model of the two tanks




Fig. 9. Labelled transition system of the controller for the two tanks model

The HYPE model is given in Figure 8. The values of r_{in} , r_{out} , s_{in} , s_{out} determine the behaviour of the system. Since the system starts with tank A draining, if $r_{in} \geq r_{out}$ then liquid will only be taken from tank A . When $r_{in} > r_{out}$ then the tank will fill repeatedly. If $r_{in} < r_{out}$ and $s_{in} \geq s_{out}$ then there will be a single switch to tank B after which all liquid will be drained by tank B . This means that the most interesting behaviour occurs in the situation when $r_{in} < r_{out}$ and $s_{in} < s_{out}$.

Figure 9 gives the labelled transition system of the controller for the two tanks model. Recall that the labelled transition system does not consider event conditions. Here $C_1 = C_A \parallel C_B \parallel C_{dA}$ and $C_2 = C_A \parallel C_B \parallel C_{dB}$. There are many cycles in this system, but for the purposes of this example, the two of interest are the infinite sequence of events that consists of $\underline{empty}_A, \underline{empty}_B$ repeated, and the infinite sequence that consists of $\underline{full}_A, \underline{full}_B, \underline{empty}_A, \underline{empty}_B$. The I-graph for this controller is given in Figure 10. It only includes vertices that are reachable from vertices where κ consists of ones. As can be seen, it is not possible to obtain the second of the sequences, since \underline{full}_A inhibits \underline{empty}_A and \underline{full}_B does not enable it. In fact, there are no cycles involving \underline{full}_A or \underline{full}_B . However, the cycle consisting of $\underline{empty}_A, \underline{empty}_B$ is possible, and it occurs when the levels of tanks A and B are both initialised to zero. Otherwise, by induction on the number of \underline{empty}_A and \underline{empty}_B that have fired, it can be shown that the cycle does not happen. This example shows both the limits of I-graphs but also how it is possible to detect undesirable situations.

Note that by having a separate controller, we are able to reason about the behaviour of the model, and either show that certain bad behaviours do not happen, or at least alert the modeller to the possibilities of these behaviours, which then may be ruled out by further investigation of specific parameters of the model.

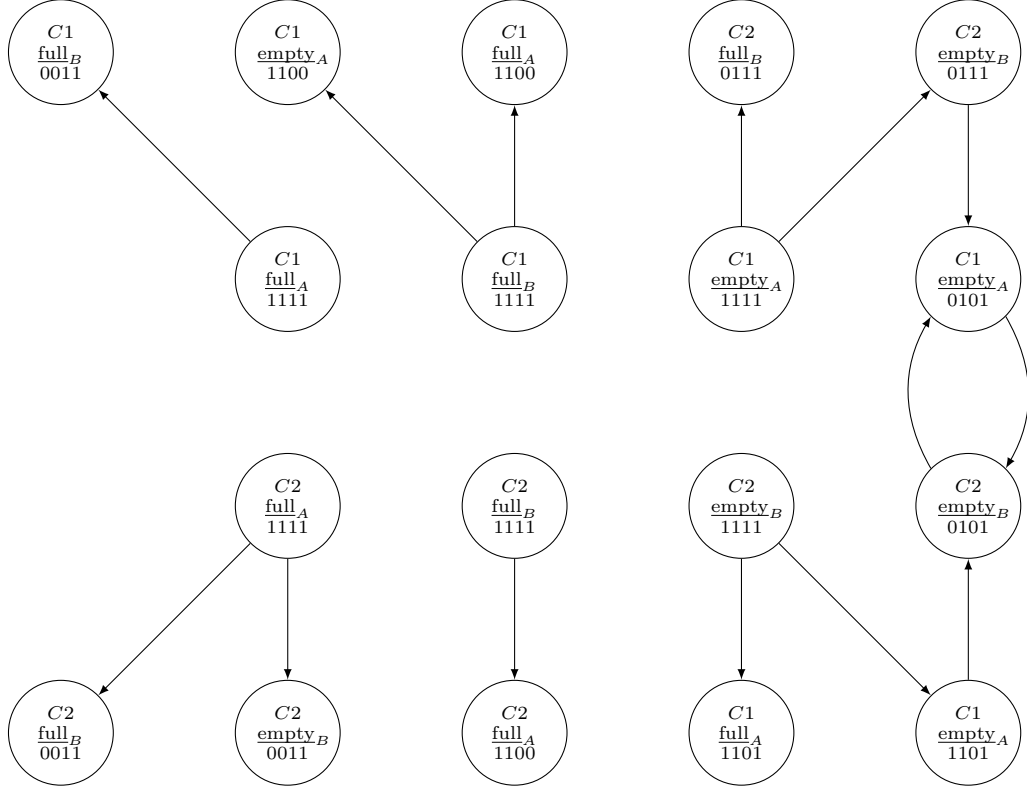


Fig. 10. I-graph for the two tanks model where the order of the κ vector is $\underline{\text{full}}_A$, $\underline{\text{empty}}_A$, $\underline{\text{full}}_B$, $\underline{\text{empty}}_B$.

6.1. Ensuring well-behavedness

We now define the property of the I-graph that ensures its HYPE model is well-behaved.

Theorem 6.1. A HYPE model with an acyclic I-graph is well-behaved.

Earlier, we defined a HYPE model as well-behaved by the absence of infinite sequences of instantaneous events. Since the detailed behaviour of a HYPE model is defined in terms of a hybrid automaton, this result must be proved in terms of hybrid automata and we introduce the appropriate definitions to achieve this. First, we define a notion of behaviour in a hybrid automaton.

Definition 6.7. Given a hybrid automaton \mathcal{H} , a sequence of simultaneous instantaneous events is a finite or infinite trace

$$(v_{i_1}, \mathbf{x}_{i_1}) \xrightarrow{\underline{a}_{i_1}} (v_{i_2}, \mathbf{x}_{i_2}) \xrightarrow{\underline{a}_{i_2}} (v_{i_3}, \mathbf{x}_{i_3}) \xrightarrow{\underline{a}_{i_3}} \dots$$

such that the events happen in the same time instance, $\text{act}(\underline{a}_{i_k})$ is true for \mathbf{x}_{i_k} and $\text{res}(\underline{a}_{i_k})$ maps \mathbf{x}_{i_k} to $\mathbf{x}_{i_{k+1}}$

This implies that $\text{act}(\underline{a}_{i_k}) \neq \perp$ for all i_k which is correct because we are only considering events that happen in the same time point rather than after some unspecified delay. From this definition, we can prove the following lemma.

Lemma 6.1. Given a HYPE model P , let $(v_i, \mathbf{x}_i) \xrightarrow{\underline{a}_i} (v_j, \mathbf{x}_j) \xrightarrow{\underline{a}_j}$ be a sequence of simultaneous instantaneous transitions in $\mathcal{H}(P)$ then $\mathbf{x}_j \in R_{\underline{a}_i} \cap G_{\underline{a}_j}$.

Proof. By definition of the behaviour of hybrid automata. □

Definition 6.8. Given a well-defined HYPE model $P \stackrel{\text{def}}{=} \Sigma \bowtie_* \text{init}.Con$, and a configuration $v = \langle \Sigma \bowtie_* C, \sigma \rangle$ in $\mathcal{H}(P)$, $con(v) = C$. For an automaton state (v, \mathbf{x}) , $con((v, \mathbf{x})) = con(v)$.

Since P is well-defined, the model must have the form $\langle \Sigma \bowtie_* C, \sigma \rangle$ and additionally, $C \in \text{ds}(\text{init}.Con)$, both from Proposition 5.2.

We now have the following results about I-graphs relating them to the behaviour of a hybrid automaton.

Lemma 6.2. Given a HYPE model P , let $(v_{i_1}, \mathbf{x}_{i_1}) \xrightarrow{a_{i_1}} (v_{i_2}, \mathbf{x}_{i_2}) \xrightarrow{a_{i_2}} (v_{i_3}, \mathbf{x}_{i_3}) \xrightarrow{a_{i_3}} \dots$ be an infinite or finite sequence of instantaneous events firing in the same time instant in $\mathcal{H}(P)$. Then

$$(con(v_{i_1}), \underline{a}_{i_1}, 1^h) \rightarrow (con(v_{i_2}), \underline{a}_{i_2}, \kappa_{i_2}) \rightarrow (con(v_{i_3}), \underline{a}_{i_3}, \kappa_{i_3}) \rightarrow \dots$$

is a path in G , the I-graph of P where $\kappa_{i_{k+1}}$ is constructed from κ_{i_k} as described in Definition 6.6.

Proof. First we consider the induction step. Suppose $(v_{i_{k-1}}, \mathbf{x}_{i_{k-1}}) \xrightarrow{a_{i_{k-1}}} (v_{i_k}, \mathbf{x}_{i_k}) \xrightarrow{a_{i_k}} (v_{i_{k+1}}, \mathbf{x}_{i_{k+1}}) \xrightarrow{a_{i_{k+1}}}$ and we know that $(con(v_{i_{k-1}}), \underline{a}_{i_{k-1}}, \kappa_{i_{k-1}}) \rightarrow (con(v_{i_k}), \underline{a}_{i_k}, \kappa_{i_k})$. We need to show that the I-graph contains the edge $(con(v_{i_k}), \underline{a}_{i_k}, \kappa_{i_k}) \rightarrow (con(v_{i_{k+1}}), \underline{a}_{i_{k+1}}, \kappa_{i_{k+1}})$. The transition $(v_{i_k}, \mathbf{x}_{i_k}) \xrightarrow{a_{i_k}} (v_{i_{k+1}}, \mathbf{x}_{i_{k+1}})$ is only possible if there is a transition in the labelled transition system of P , and moreover, this is only possible if there is such a transition in the controller, so $\langle con(v_{i_k}), \sigma \rangle \xrightarrow{a_{i_k}} \langle con(v_{i_{k+1}}), \sigma \rangle$.

Since $(v_{i_k}, \mathbf{x}_{i_k}) \xrightarrow{a_{i_k}} (v_{i_{k+1}}, \mathbf{x}_{i_{k+1}}) \xrightarrow{a_{i_{k+1}}}$ and by Lemma 6.1, \underline{a}_{i_k} must have been enabled hence $\kappa_{i_k}[i_k] = 1$. Similarly, since $(v_{i_{k+1}}, \mathbf{x}_{i_{k+1}}) \xrightarrow{a_{i_{k+1}}}$, we know $\kappa_{i_{k+1}}[i_{k+1}] = 1$. Finally, κ_{i_k} is correct by induction and $\kappa_{i_{k+1}}$ is an update of κ_{i_k} taking into account which events a_{i_k} has enabled and inhibited. Since there is a node with $C_{i_{k+1}}$ and $\underline{a}_{i_{k+1}}$ for each possible κ , the edge can be constructed. Hence the four conditions to be satisfied for an edge in the I-graph have been satisfied.

Returning to the base case, the same argument can be used for the first condition and third condition, the second condition is immediate and the fourth condition follows since κ_{i_2} captures which events are inhibited by \underline{a}_{i_1} . \square

This lemma allows us to now prove the theorem stated at the start of this section, and thereby characterise the I-graphs that give us well-behaved HYPE models.

Proof of Theorem 6.1: For a HYPE model P , we need to show that $\mathcal{H}(P)$ has a finite number of finite sequences of instantaneous events that can fire in the same time instance and these sequences are independent of the initial state of the system.

From Lemma 6.2, the I-graph of a HYPE model captures the sequences of instantaneous events which can occur in the execution of $\mathcal{H}(P)$ since by construction, all events appear in the controller. If there are no cycles in the I-graph, which is finite by construction, then there are no infinite sequences possible in the execution. Since an I-graph is finite, there can only be a finite number of finite sequences in $\mathcal{H}(P)$. Finally, since there are nodes in the I-graph where the vector κ consist of 1's which means that all events are active, this graph is independent of the initial state. \square

6.2. Results for well-behaved HYPE models

We have now defined well-behavedness for HYPE models and shown how the I-graph of the model can be used to check for this behaviour. Additionally, this has been related to the hybrid semantics of the models. We now consider the implications of this definition and result and how we can ensure that it is a decidable property. As mentioned above, it is not always practicable to determine $G_{\underline{a}}$ or $R_{\underline{a}}$. We can extend the definition of an I-graph as follows.

Definition 6.9. Consider a HYPE model with $R_{\underline{a}} = r_{\underline{a}}(G_{\underline{a}})$ for each event \underline{a} , and let $R'_{\underline{a}}, G'_{\underline{a}} \subseteq \mathbb{R}^n$, be a collection of sets such that for all \underline{a} , $R_{\underline{a}} \subseteq R'_{\underline{a}}$ and $G_{\underline{a}} \subseteq G'_{\underline{a}}$. These sets are called *extended sets*.

Definition 6.10. An *extended I-graph* of a HYPE model is defined in the same manner as its I-graph but where extended sets are used to determine enablers and inhibitors.

Our result from the previous section can be carried over to extended sets and hence if the exact sets are not computable or very hard to compute, we can work with an overapproximation.

Theorem 6.2. A HYPE model with an acyclic extended I-graph is well-behaved.

Proof. By increasing the size of the sets of values for activation conditions and resets, additional edges may be added to the I-graph but no edge can be removed. Hence if the extended I-graph is acyclic, so is the I-graph. \square

Corollary 6.1. Checking whether a HYPE model with computable extended sets has an acyclic I-graph is decidable.

Proof. Since an I-graph is finite, and the sets can be computed, it can be effectively constructed and checked for acyclicity. \square

If we know the controller has a specific cyclic form, we can also investigate conditions that will prevent a cycle from occurring in the I-graph. These results apply in the case of extended sets as well.

Proposition 6.1. Let P be a HYPE model with $Con \stackrel{def}{=} \underline{a}_1 \dots \underline{a}_n \cdot Con$. If the resets for all \underline{a}_i are the identity and $\bigcap_{\underline{a}_i \in Con} G_{\underline{a}_i} = \emptyset$ then P is well-behaved. It is straightforward to see that $R_{\underline{a}} = G_{\underline{a}}$ and at some point there is no overlap between the sets of two different events.

Proof. The only possible cycle is that consisting of the events $\underline{a}_1, \dots, \underline{a}_n$ repeated. Moreover, $\bigcap_{\underline{a}_i \in Con} G_{\underline{a}_i} = \emptyset$ implies that there exist j and k such that $G_{\underline{a}_j} \cap G_{\underline{a}_k} = \emptyset$. Additionally $R_{\underline{a}_i} = G_{\underline{a}_i}$ for all i . Hence $R_{\underline{a}_j} \cap G_{\underline{a}_k} = \emptyset$ and $R_{\underline{a}_j} \cap G_{\underline{a}_k} = \emptyset$ and hence these events inhibit each other and prevent the formation of a cycle in the I-graph. \square

Proposition 6.2. Let P be a HYPE model with $Con \stackrel{def}{=} \underline{a}_1 \dots \underline{a}_n \cdot Con$. If there exists i such that $R_{\underline{a}_i} \cap G_{\underline{a}_{i+1}}$ is empty (where addition is modulo n) then P is well-behaved.

Proof. The lack of overlap of values at one point in the sequence of the controller prevents a cycle in the I-graph. \square

Now that we have presented the basic theory of well-definedness, we can consider how this can be applied compositionally. We prove the results for the collections of sets $G_{\underline{a}}$ and $R_{\underline{a}}$, but since none of the results depend on the set being minimal, they also apply to extended sets. The first result is to be expected since the events of the two systems cannot affect each other.

Proposition 6.3. Let Con and Con' be two controllers such that all events in Con are pairwise independent from those of Con' . If Con and Con' are well-behaved then $Con \parallel Con'$ is well-behaved.

Proof. The independence of events implies that $\text{ev}(Con) \cap \text{ev}(Con') = \emptyset$, hence the use of \parallel . Let the G be the I-graph of $Con \boxtimes Con'$. Its vertices are in the set $ds(\text{init.}(Con \boxtimes Con')) \times (\text{ev}(Con) \cup \text{ev}(Con')) \times \{0, 1\}^h$. Assume there is an cycle in G . We will show it is possible to construct a cycle in H , the I-graph of Con .

Let $(C_1 \boxtimes_* C'_1, \underline{a}_1, \kappa_1) \rightarrow (C_2 \boxtimes_* C'_2, \underline{a}_2, \kappa_2) \rightarrow (C_3 \boxtimes_* C'_3, \underline{a}_3, \kappa_3)$ be a subpath of the cycle in G such that $\underline{a}_3 \in \text{ev}(Con)$ and $\underline{a}_2 \in \text{ev}(Con')$. Since the events of the two controllers are disjoint, $C_2 = C_3$. Since the events of Con and Con' are pairwise independent, \underline{a}_2 can have no effect on the activation of \underline{a}_3 and hence $\kappa_2[3]$ must be 1.

Therefore, we know there is an edge $(C_1 \boxtimes_* C'_1, \underline{a}_1, \kappa_1) \rightarrow (C_3 \boxtimes_* C'_3, \underline{a}_3, \kappa_3)$ and $(C_2 \boxtimes_* C'_2, \underline{a}_2, \kappa_2)$ can be removed from the cycle, while ensuring a cycle remains. Similarly all other events from Con' can be removed from the cycle and the remaining cycle will consist only of events from Con . If we modify the vertices of G by changing any $C_i \boxtimes_* C'_i$ to C_i and appropriately removing elements from the vectors, we obtain vertices from H . Moreover, this cycle must appear in H since it is constructed purely with reference to events in Con . Thus, we have a contradiction. \square

Proposition 6.4. Let Con and Con' be two controllers such that for all $\underline{a} \in \text{ev}(Con) \setminus \text{ev}(Con')$ and for $\underline{a}' \in \text{ev}(Con') \setminus \text{ev}(Con)$, no \underline{a} activates an \underline{a}' and no \underline{a}' activates an \underline{a} . If Con and Con' are well-behaved then $Con \boxtimes_* Con'$ is well-behaved.

Proof. Let the G be the I-graph of $Con \boxtimes_* Con'$ with vertices $ds(\text{init.}(Con \boxtimes_* Con')) \times (\text{ev}(Con) \cup \text{ev}(Con')) \times \{0, 1\}^h$. Assume there is a cycle in G . We will show this implies that there is a cycle in H , the I-graph of Con .

Let $(C_1 \boxtimes_* C'_1, \underline{a}_1, \kappa_1) \rightarrow (C_2 \boxtimes_* C'_2, \underline{a}_2, \kappa_2) \rightarrow (C_3 \boxtimes_* C'_3, \underline{a}_3, \kappa_3)$ be a subpath of the cycle in G . First consider the case where $\underline{a}_3 \in \text{ev}(Con)$ and $\underline{a}_2 \in \text{ev}(Con')$ and $\underline{a}_2, \underline{a}_3 \notin \text{ev}(Con) \cap \text{ev}(Con')$. Since the events are in separate controllers, $C_2 = C_3$. Since the unshared events of Con and Con' cannot activate each other, \underline{a}_2 can have no effect on the activation of \underline{a}_3 and hence $\kappa_2[3]$ must be 1.

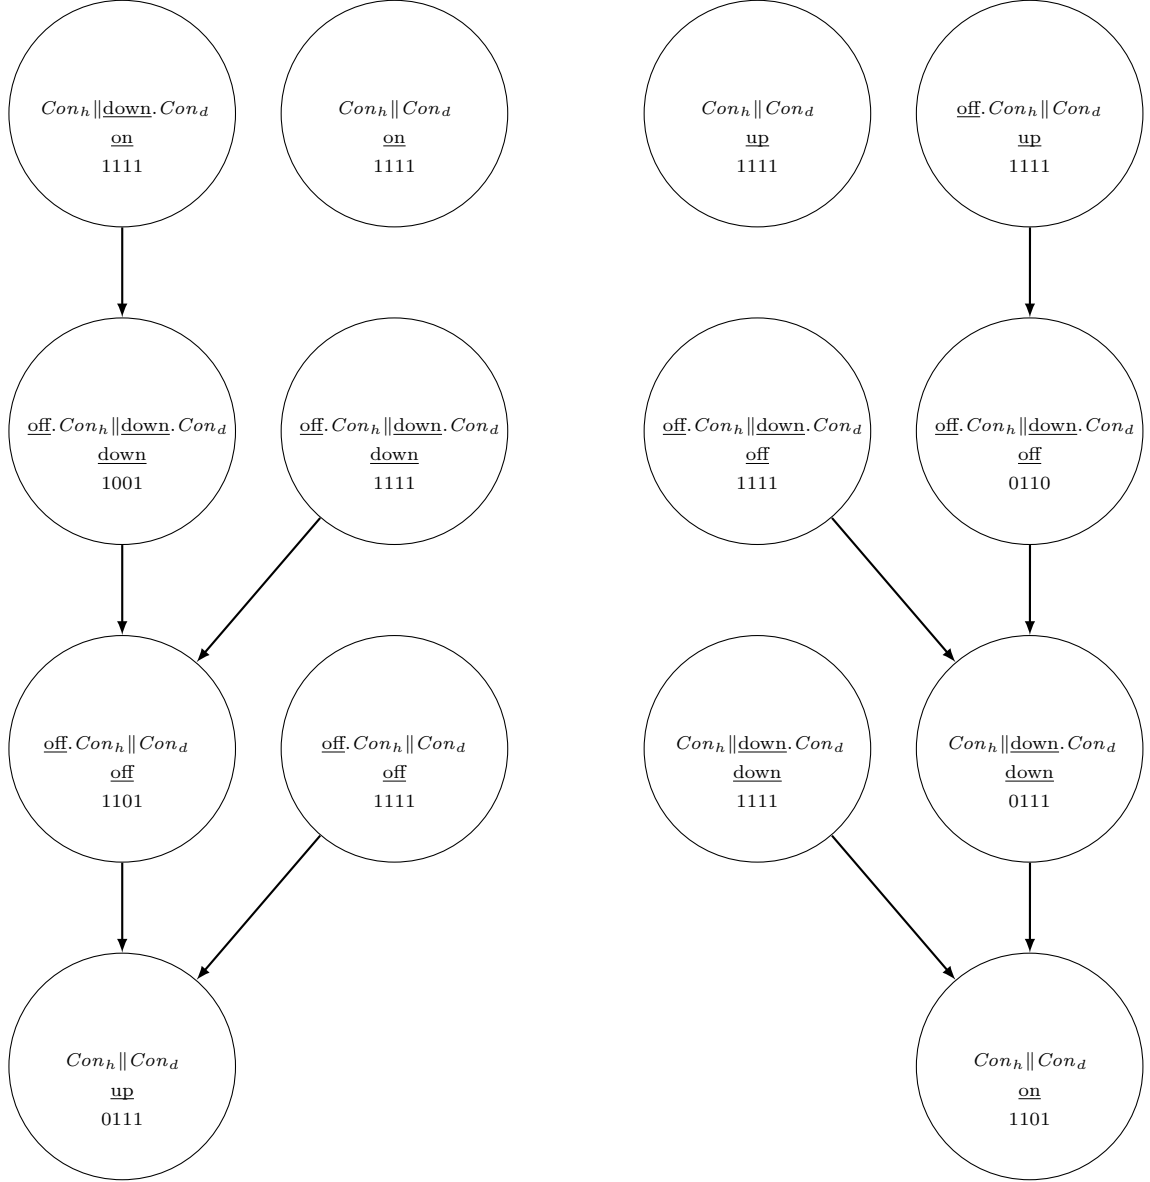


Fig. 11. I-graph for $Con_h \parallel Con_d$ where the order of the κ vector is on, off, up, down.

Therefore, we know there is an edge $(C_1 \boxtimes_* C'_1, \underline{a}_1, \kappa_1) \rightarrow (C_3 \boxtimes_* C'_3, \underline{a}_3, \kappa_3)$ and $(C_2 \boxtimes_* C'_2, \underline{b}_2, \kappa_2)$ can be removed from the cycle, while ensuring a cycle remains. Similarly all other events from Con' can be removed from the cycle and the remaining cycle will consist only of unshared events from Con_1 , and events that occur in both controllers.

Next let $(C_1 \boxtimes_* C'_1, \underline{a}_1, \kappa_1) \rightarrow (C_2 \boxtimes_* C'_2, \underline{a}_2, \kappa_2) \rightarrow (C_3 \boxtimes_* C'_3, \underline{a}_3, \kappa_3)$ be a subpath of this reduced cycle in G with $\underline{a}_1, \underline{a}_3 \in \text{ev}(Con) \setminus \text{ev}(Con')$ and $\underline{a}_2 \in \text{ev}(Con) \cap \text{ev}(Con')$. By the definition of I-graphs, there must be a transition $\langle C_2 \boxtimes_* C'_2, \sigma \rangle \xrightarrow{\underline{a}_2} \langle C_3 \boxtimes_* C'_3, \sigma \rangle$ in the labelled transition system of $Con \boxtimes_* Con'$, hence we know that the transition $\langle C_2, \sigma \rangle \xrightarrow{\underline{a}_2} \langle C_3, \sigma \rangle$ exists, because of the use of \boxtimes_* . If we can show that this transition leads to an edge in the graph, then by the same reasoning in the previous proof, we can conclude that we have a cycle in H .

We consider a path $(C_1, \underline{a}_1, \kappa'_1) \rightarrow (C_2, \underline{a}_2, \kappa'_2) \rightarrow (C_3, \underline{a}_3, \kappa'_3)$. We require $\kappa'_3[3] = 1$ to ensure the existence

of the second edge, regardless of the value of $\kappa'_1[3]$ and $\kappa'_2[3]$. If $\kappa'_1[3] = 1$, then \underline{a}_1 and \underline{a}_2 must not inhibit \underline{a}_3 , and if $\kappa'_1[3] = 0$, \underline{a}_2 must activate \underline{a}_3 . But we know this already from the path in the reduced cycle in G . Hence we can conclude we have a cycle in H and we have shown a contradiction. \square

Finally, in this section, we consider the two examples. For the two tanks example, we can apply Proposition 6.3 to determine that the system with $Con_A \parallel Con_B$ is well-behaved. Since the I-graph of Con_{AB} contains a cycle, we cannot go further than this.

Considering the orbiter, if $k_2 < k_1$, then Con_h has well-behaved resets by Proposition 6.1 and similarly, if $k_4 < k_3$, then Con_d has well-behaved resets. However, neither Proposition 6.3 nor Proposition 6.4 apply because the events are not independent and activation is possible between events in each controller. Hence, the I-graph of $Con_h \parallel Con_d$ must be constructed using the fact that $k_1 < k_4$ which allows us to determine the enabling and inhibition of events by the other events. Figure 11 shows the I-graph, which is acyclic. Additionally, Con_s has well-behaved resets by Proposition 6.2. Hence, using Proposition 6.3, the system with the controller $(Con_h \bowtie_* Con_d) \bowtie_* Con_s$ is well-behaved.

The condition of having an acyclic I-graph is a reasonable combination of practical expressiveness and effective computability. It allows us to determine which controllers surely do not have instantaneous Zeno behaviour through a simple graph construction. Moreover, by choosing computable sets, we can ensure decidability of the process. In our experience, for reasonable models, the acyclicity condition should not yield too many false positives. Also, I-graphs are unlikely to be too large as only few vertices can be reached from those where κ is the vector consisting of ones and construction can be done on the fly.

Note that the I-graph construction does not take into account the initial values of variables. The fact that κ is initially set to ones in the construction means that all events are initially possible. This has the advantage that the I-graph applies to all behaviours of the HYPE model, regardless of initial values but the disadvantage that this leads to an overapproximation of behaviour. An alternative approach would be to construct a graph from the initial conditions, but this would require a more complex algorithm with difficult termination conditions. This would have been more precise in the sense of being a smaller overapproximation; however, the cost does not appear to be justified by the benefits. As we can see from the two examples, the approach taken here, is sufficient to reason about these controllers.

To sum up, the separate definition of the controller in the HYPE model, together with the event conditions allow us to reason about the behaviour of the model without considering its detailed behaviour in terms of a hybrid automaton. This supports the design of models that exclude unwanted behaviour.

7. Comparison with hybrid automata

We now compare the expressive power of HYPE with that of hybrid automata at the level of composition of models. We first need to describe how two different HYPE systems can be composed in a synchronisation. The explanation of the choices made are given after the definition.

Definition 7.1. Let $P_i \stackrel{\text{def}}{=} \Sigma_i \bowtie_* \text{init}.Con_i$ for $i = 1, 2$ be two well-defined HYPE systems given by $(P_i, \mathcal{V}_i, IN_i, IT_i, \mathcal{E}_i, \mathcal{A}_i, ec_i, iv_i, EC_i, ID_i)$ respectively. The *synchronisation* of P_1 and P_2 ($P_1 \otimes P_2$) is the HYPE model $(P, \mathcal{V}_1 \cup \mathcal{V}_2, IN_1 \cup IN_2, IT_1 \cup IT_2, \mathcal{E}_1 \cup \mathcal{E}_2, \mathcal{A}_1 \cup \mathcal{A}_2, ec, iv_1 \cup iv_2, EC_1 \cup EC_2, ID_1 \cup ID_2)$ which is defined whenever $IN_1 \cap IN_2 = \emptyset$ and $IT_1 \cap IT_2 = \emptyset$ and where the well-defined system is $P \stackrel{\text{def}}{=} (\Sigma_1 \bowtie_* \Sigma_2) \bowtie_* \text{init}.(Con_1 \bowtie_* Con_2)$ and ec is defined by

$$ec(\underline{a}) = \begin{cases} ec_1(\underline{a}) & \text{if } \underline{a} \notin \mathcal{E}_2 \\ ec_2(\underline{a}) & \text{if } \underline{a} \notin \mathcal{E}_1 \\ (act_1(\underline{a}) \wedge act_2(\underline{a}), res_1(\underline{a}) \wedge res_2(\underline{a})) & \text{if } \underline{a} \in \mathcal{E}_1 \cap \mathcal{E}_2 \end{cases}$$

The condition $IN_1 \cap IN_2 = \emptyset$ requires influence names to be disjoint. This is necessary to ensure a well-defined synchronisation. Note that this condition implies that $\mathcal{A}_1 \cap \mathcal{A}_2 = \emptyset$ and $iv_1 \cap iv_2 = \emptyset$. In the case where the same influence name is used in two different models, one can be renamed to a fresh name, as they are essentially arbitrary names used to identify flows (unlike variables which have specific meanings). After renaming, the synchronisation is well-defined.

The condition $IT_1 \cap IT_2 = \emptyset$ requires that influence types are disjoint. It would be possible to allow shared influence types with identical definitions up to substitution of variable names based on a bijection between the variables, but our choice is to go for the simpler definition.

Once these conditions hold, the union of most of the sets can be taken straightforwardly. This allows for influences from the different models to map to a shared variable which is an important aspect of the synchronisation definition. For the definition of ec , we need to define $\perp \wedge \phi = \phi \wedge \perp = \phi$ and $\perp \wedge \perp = \perp$. This means that for a shared event, it is only non-urgent if it is non-urgent in both components. It is reasonable to treat urgency as stronger than non-urgency.

Proposition 7.1. Given two well-defined HYPE models P_1 and P_2 , $P_1 \otimes P_2$ is well-defined.

Proof. The first condition about the form of subcomponents holds because it holds in both models. The second condition about disjointness of influences between subcomponents holds because influences are disjoint between models. The conditions requiring cooperation on all shared events hold by the definition of the combined controlled systems. Since the two models are well-defined, $\text{ev}(\Sigma_i) = \mathcal{E}_i$, $\text{in}(\Sigma_i) = IN_i$, $\text{ev}(Con_i) \subseteq \text{ev}(\Sigma_i)$ for $i = 1, 2$, and hence we know that $\text{ev}(\Sigma_1 \boxtimes \Sigma_2) = \mathcal{E}_1 \cup \mathcal{E}_2$, $\text{in}(\Sigma_1 \boxtimes \Sigma_2) = IN_1 \cup IN_2$ and $\text{ev}(Con_1 \boxtimes Con_2) \subseteq \text{ev}(\Sigma_1 \boxtimes \Sigma_2)$. \square

To compare synchronisation of HYPE models with a similar construction on hybrid automata, we work with the definition of product as given in [HH94] modified to take account of urgent actions.

Definition 7.2. Let H_1 and H_2 be two hybrid automata with $\mathcal{H}_i = (V_i, E_i, \mathbf{X}_i, \mathcal{E}_i, \text{flow}_i, \text{init}_i, \text{inv}_i, \text{event}_i, \text{jump}_i, \text{reset}_i, \text{urgent}_i)$ for $i = 1, 2$. The *synchronised product* (or *parallel composition*) of H_1 and H_2 ($H_1 \times H_2$) is $H = (V_1 \times V_2, E, \mathbf{X}_1 \cup \mathbf{X}_2, \mathcal{E}_1 \cup \mathcal{E}_2, \text{flow}, \text{init}, \text{inv}, \text{event}, \text{jump}, \text{reset}, \text{urgent})$ where

- $(e_1, e_2) = ((v_1, v_2), (v'_1, v'_2)) \in E$ if
 1. $v_2 = v'_2$ and $\text{event}_1(e_1) \notin \mathcal{E}_2$ or
 2. $v_1 = v'_1$ and $\text{event}_2(e_2) \notin \mathcal{E}_1$ or
 3. $\text{event}_1(e_1) = \text{event}_2(e_2) \in \mathcal{E}_1 \cap \mathcal{E}_2$
- if $e = (e_1, e_2)$ with
 1. $\text{event}_1(e_1) \notin \mathcal{E}_2$ then $\text{event}(e) = \text{event}_1(e_1)$, $\text{jump}(e) = \text{jump}_1(e_1)$, $\text{reset}(e) = \text{reset}_1(e_1)$, $\text{urgent}(e) = \text{urgent}_1(e_1)$
 2. $\text{event}_2(e_2) \notin \mathcal{E}_1$ then $\text{event}(e) = \text{event}_2(e_2)$, $\text{jump}(e) = \text{jump}_2(e_2)$, $\text{reset}(e) = \text{reset}_2(e_2)$, $\text{urgent}(e) = \text{urgent}_2(e_2)$
 3. $\text{event}_1(e_1) = \text{event}_2(e_2) \in \mathcal{E}_1 \cap \mathcal{E}_2$ then $\text{event}(e) = \text{event}_1(e_1)$, $\text{jump}(e) = \text{jump}_1(e_1) \wedge \text{jump}_2(e_2)$, $\text{reset}(e) = \text{reset}_1(e_1) \wedge \text{reset}_2(e_2)$, $\text{urgent}(e) = \text{urgent}_1(e_1) \vee \text{urgent}_2(e_2)$,
- if $v = (v_1, v_2)$ then $\text{flow}(v) = \text{flow}_1(v_1) \wedge \text{flow}_2(v_2)$, $\text{init}(v) = \text{init}_1(v_1) \wedge \text{init}_2(v_2)$, $\text{inv}(v) = \text{inv}_1(v_1) \wedge \text{inv}_2(v_2)$.

For this definition to give a meaningful product, whenever $X \in \mathbf{X}_1 \cap \mathbf{X}_2$, we require that $\text{flow}_1(v_1)[X] = \text{flow}_2(v_2)[X]$ for a vertex (v_1, v_2) . If this is not the case, flow will evaluate to false for X . For consistency with the definition of synchronisation, we consider urgency stronger than non-urgency and hence $\text{urgent}(e) = \text{urgent}_1(e_2) \vee \text{urgent}_2(e_2)$.

Now that we have a definition of product for HYPE models and synchronised product for hybrid automata, we can compare them and identify whether they have the same or different expressive power.

Theorem 7.1. Let P_1 and P_2 be well-defined HYPE systems $(P_i, \mathcal{V}_i, IN_i, IT_i, \mathcal{E}_i, \mathcal{A}_i, ec_i, iw_i, EC_i, ID_i)$ for $i = 1, 2$. If $\mathcal{V}_1 \cap \mathcal{V}_2 = \emptyset$ then $\mathcal{H}(P_1 \otimes P_2) = \mathcal{H}(P_1) \times \mathcal{H}(P_2)$ when vertices unreachable from the initial vertex are excluded.

Proof. See Appendix A. \square

Corollary 7.1. Let P_1 and P_2 be well-defined HYPE systems $(P_i, \mathcal{V}_i, IN_i, IT_i, \mathcal{E}_i, \mathcal{A}_i, ec_i, iw_i, EC_i, ID_i)$ for $i = 1, 2$, with hybrid automata $\mathcal{H}(P_i) = (V_i, E_i, \mathbf{X}_i, \mathcal{E}_i, \text{flow}_i, \text{init}_i, \text{inv}_i, \text{event}_i, \text{jump}_i, \text{reset}_i, \text{urgent}_i)$, synchronisation $P_1 \otimes P_2$, and synchronised product hybrid automaton $\mathcal{H}(P_1 \otimes P_2) = (V, E, \mathbf{X}, \mathcal{E}, \text{flow}, \text{init}, \text{inv}, \text{event}, \text{jump}, \text{reset}, \text{urgent})$. If $\mathbf{X}_1 \cap \mathbf{X}_2 \neq \emptyset$, there exists $X \in \mathbf{X}_1 \cap \mathbf{X}_2$ with flow not constant zero at some mode, and $\mathcal{H}(P_1) \times \mathcal{H}(P_2)$ is defined then $\mathcal{H}(P_1 \otimes P_2) \neq \mathcal{H}(P_1) \times \mathcal{H}(P_2)$.

Proof. Since $\mathcal{H}(P_1) \times \mathcal{H}(P_2)$ is defined, then for all $X \in \mathbf{X}_1 \cap \mathbf{X}_2$, and for all vertices, $v = (v_1, v_2)$ in $\mathcal{H}(P_1) \times \mathcal{H}(P_2)$, $\text{flow}_1(v_1)[X] = \text{flow}_2(v_2)[X]$. Each $\text{flow}_i(v_i)[X] = \sum \{r[I(\mathcal{W})] \mid iw(\iota) = X \text{ and } \sigma_i(\iota) =$

$(r, I(\mathcal{W}))\}$ where $v_i = \langle \Sigma_i \bowtie D_i, \sigma_i \rangle$. Without considering the details of the right hand side of the equation, we can write $flow_1(v_1)[X] = flow_2(v_2)[X] = f(\mathbf{X})$.

Using the technique of the proof of Theorem 7.1, we can conclude that $\langle (\Sigma_1 \bowtie \Sigma_2) \bowtie (D_1 \bowtie D_2), \sigma_1 \cup \sigma_2 \rangle$ is a configuration in the labelled transition system of $P_1 \otimes P_2$ and hence it is a vertex in $\mathcal{H}(P_1 \otimes P_2)$, say v' . Hence $flow(v')[X] = \sum \{r[I(\mathcal{W})] \mid iv(\iota) = X \text{ and } (\sigma_1 \cup \sigma_2)(\iota) = (r, I(\mathcal{W}))\}$. Since $IN_1 \cap IN_2 = \emptyset$, the domains of σ_1 and σ_2 are distinct. From this, we can conclude that $flow(v')[X] = 2f(\mathbf{X})$. \square

We have shown that by restricting the HYPE models to those with disjoint sets of variables, they have the same expressive power with respect to compositionality as hybrid automata. The necessity of this restriction is shown by the corollary. When two hybrid automata have shared variables, either their product is undefined or there must be a zero flow for the shared variables, restricting the compositionality of the formalism.

By contrast, in HYPE, influences from each component of the product can map to a shared variable and hence the ODE for that variable can be determined by both components. Therefore compositionality is supported fully by the synchronisation product for HYPE, and allows for useful interaction between models. From this we can conclude that HYPE is more fine-grained and more flexible in its modelling style than hybrid automata. In many cases, changes to a model can be achieved by composition rather than by rewriting the model. In further work, we will investigate how this can provide gains in terms of analysis.

8. Equivalence Semantics

An important technique in the realm of process algebras is that of semantic equivalences, where we are able to define notions that capture the idea of same behaviour. An advantage of working with a process algebra as a language to describe systems is that we aim to understand how we can identify models that have the same behaviour at the semantic level but differ syntactically. Additionally, if we can prove congruence, we are able to substitute one equivalent component or subcomponent for another, in a model. One use of substitution is to replace a component with a smaller but equivalent component to reduce the size of the labelled transition system and the resulting hybrid automaton.

We define equivalent behaviour over the labelled transition system given in Section 3 thereby focussing on the configuration of the system rather than evaluations of continuous variables, allowing us to abstract away from these details. Since our operational semantics apply to more general terms than those specified in our HYPE models, we define our equivalence over those more general terms. Let \mathcal{C} be the set of terms obtained by the grammar $S ::= \underline{a} : \alpha.S \mid \underline{a}.S \mid S' + S'$ and $P ::= C \mid P \bowtie P$ where $C \stackrel{def}{=} S$ defines a constant. Note that $\mathcal{C}_{sys} \subseteq \mathcal{C}$ and hence we can use compositionality results such as congruence to reason about HYPE models.

We base our equivalence on bisimulation [Mil89], because it is an equivalence that makes strong distinctions relating to issues such as deadlock and branching whereas language/trace equivalence cannot make these distinctions [vG90].

Definition 8.1. A relation $B \subseteq \mathcal{C} \times \mathcal{C}$ is a *system bisimulation* if for all $(P, Q) \in B$, for all $\underline{a} \in \mathcal{E}$, for all $\sigma \in \mathcal{S}$ whenever

1. $\langle P, \sigma \rangle \xrightarrow{\underline{a}} \langle P', \sigma' \rangle$, there exists $\langle Q', \sigma' \rangle$ with $\langle Q, \sigma \rangle \xrightarrow{\underline{a}} \langle Q', \sigma' \rangle$, $(P', Q') \in B$.
2. $\langle Q, \sigma \rangle \xrightarrow{\underline{a}} \langle Q', \sigma' \rangle$, there exists $\langle P', \sigma' \rangle$ with $\langle P, \sigma \rangle \xrightarrow{\underline{a}} \langle P', \sigma' \rangle$, $(P', Q') \in B$.

P and Q are *system bisimilar*, $P \sim_s Q$ if they are in a system bisimulation.

It is straightforward to show that system bisimulation is a congruence for our operators.

Theorem 8.1. \sim_s is a congruence for Prefix, Choice and Cooperation.

Proof. Let $P_1 \sim_s P_2$. Symmetrical cases have been omitted.

Prefix with influence We have the transition $\langle \underline{a} : (\iota, r, I).P_1, \sigma \rangle \xrightarrow{\underline{a}} \langle P_1, \sigma[\iota \mapsto (r, I)] \rangle$ and likewise the transition $\langle \underline{a} : (\iota, r, I).P_2, \sigma \rangle \xrightarrow{\underline{a}} \langle P_2, \sigma[\iota \mapsto (r, I)] \rangle$. Since $P_1 \sim_s P_2$, then we can conclude that $\underline{a} : (\iota, r, I).P_1 \sim_s \underline{a} : (\iota, r, I).P_2$.

Prefix without influence Similar to the previous case.

$$\begin{aligned}
D_1 &\stackrel{\text{def}}{=} \underline{\text{full}}_A.D_1 + \underline{\text{full}}_B.D_1 + \underline{\text{empty}}_A.D_2 & D_2 &\stackrel{\text{def}}{=} \underline{\text{full}}_A.D_2 + \underline{\text{full}}_B.D_2 + \underline{\text{empty}}_B.D_1 \\
S &\stackrel{\text{def}}{=} ((In_A \bowtie_* Out_A) \bowtie_* (In_B \bowtie_* Out_B)) \bowtie_* \underline{\text{init}}.D_1
\end{aligned}$$

Fig. 12. Alternative controller for two tanks

Choice There are two cases. First if $\langle P_1, \sigma \rangle \xrightarrow{\underline{a}} \langle P'_1, \sigma' \rangle$, then we have $\langle P_1 + Q, \sigma \rangle \xrightarrow{\underline{a}} \langle P'_1, \sigma' \rangle$ and since $P_1 \sim_s P_2$, $\langle P_2, \sigma \rangle \xrightarrow{\underline{a}} \langle P'_2, \sigma' \rangle$ with $P'_1 \sim_s P'_2$, and hence $\langle P_2 + Q, \sigma \rangle \xrightarrow{\underline{a}} \langle P'_2, \sigma' \rangle$ as required. Second, $\langle Q, \sigma \rangle \xrightarrow{\underline{a}} \langle Q', \sigma'' \rangle$ and $\langle P_1 + Q, \sigma \rangle \xrightarrow{\underline{a}} \langle Q', \sigma'' \rangle$. Also $\langle P_2 + Q, \sigma \rangle \xrightarrow{\underline{a}} \langle Q', \sigma'' \rangle$ and we know $Q' \sim_s Q'$.

Cooperation We need to show that $B = \{(P_1 \bowtie_L Q, P_2 \bowtie_L Q) \mid P_1 \sim_s P_2\}$ is a system bisimulation. There are three cases. First, if $\langle P_1, \sigma \rangle \xrightarrow{\underline{a}} \langle P'_1, \sigma' \rangle$ with $\underline{a} \notin L$ then $\langle P_1 \bowtie_L Q, \sigma \rangle \xrightarrow{\underline{a}} \langle P'_1 \bowtie_L Q, \sigma' \rangle$. Since $P_1 \sim_s P_2$, $\langle P_2, \sigma \rangle \xrightarrow{\underline{a}} \langle P'_2, \sigma' \rangle$ with $P'_1 \sim_s P'_2$, and hence $\langle P_2 \bowtie_L Q, \sigma \rangle \xrightarrow{\underline{a}} \langle P'_2 \bowtie_L Q, \sigma' \rangle$ with $(P'_1 \bowtie_L Q, P'_2 \bowtie_L Q) \in B$ as required. Second, if $\langle Q, \sigma \rangle \xrightarrow{\underline{a}} \langle Q', \sigma'' \rangle$ with $\underline{a} \notin L$ then the proof is straightforward and clearly $(P_1 \bowtie_L Q', P_2 \bowtie_L Q') \in B$. Third and lastly, $\underline{a} \in L$ and $\langle P_1, \sigma \rangle \xrightarrow{\underline{a}} \langle P'_1, \sigma' \rangle$ and $\langle Q, \sigma \rangle \xrightarrow{\underline{a}} \langle Q', \sigma'' \rangle$, then $\langle P_1 \bowtie_L Q, \sigma \rangle \xrightarrow{\underline{a}} \langle P'_1 \bowtie_L Q', \Gamma(\sigma, \sigma', \sigma'') \rangle$. Since $P_1 \sim_s P_2$, $\langle P_2, \sigma \rangle \xrightarrow{\underline{a}} \langle P'_2, \sigma' \rangle$ with $P'_1 \sim_s P'_2$, and hence $\langle P_2 \bowtie_L Q, \sigma \rangle \xrightarrow{\underline{a}} \langle P'_2 \bowtie_L Q', \Gamma(\sigma, \sigma', \sigma'') \rangle$. $(P'_1 \bowtie_L Q', P'_2 \bowtie_L Q') \in B$ as required. \square

It can be remarked that \sim_s is a *robust* or *stateless bisimulation* [CR03, MRG05]. This style of equivalence, which relates process terms for arbitrary states or data, has been used extensively in hybrid process algebras. Moreover, our transition system specification in Figure 2 conforms to the *process-tyft* format which is a congruence format for stateless bisimulation and this offers an alternative route to proving the theorem above.

Given bisimilar controllers, we can show that an uncontrolled system combined with either will result in the same behaviour in terms of system bisimulation.

Corollary 8.1. Given an uncontrolled systems Σ and two controllers such that $Con_1 \sim_s Con_2$ then $\Sigma \bowtie_L \underline{\text{init}}.Con_1 \sim_s \Sigma \bowtie_L \underline{\text{init}}.Con_2$ if these controlled systems are well-defined.

Proof. By congruence. \square

Consider the alternative controller in Figure 12 for the two tanks of Section 6. Controllers $C_A \parallel C_B \parallel C_{dA}$ and D_1 are clearly isomorphic as can be seen from Figure 9, and hence are system bisimilar, because no state changes can happen during the execution of the controllers on their own. We can apply the corollary and conclude that *System* and S are system bisimilar.

The preceding congruence theorem is useful when reasoning about a single model and proving later results about bisimilar systems; however, due to the constrained form of HYPE models, we are also interested in congruence with respect to the synchronisation operator between HYPE models, and we now prove this. We need to lift the definition of system bisimulation to HYPE models, and this is achieved by requiring all elements of the tuple except the first are identical.

Definition 8.2. Given two well-defined HYPE models $(P_i, \mathcal{V}, IN, IT, \mathcal{E}, \mathcal{A}, ec, iv, EC, ID)$ for $i = 1, 2$, they are *system bisimilar as models* ($P_1 \sim_{\text{sm}} P_2$) if $P_1 \sim_s P_2$.

Theorem 8.2. Let P_1, P_2 and Q be well-defined HYPE models. If $P_1 \sim_{\text{sm}} P_2$ then $P_1 \otimes Q \sim_{\text{sm}} P_2 \otimes Q$ and $Q \otimes P_1 \sim_{\text{sm}} Q \otimes P_2$.

Proof. First, note that the synchronisations are defined because P_1 and P_2 have the same elements in their tuples (excluding the first element). This means that the synchronisations will also have this property and hence can be system bisimilar as models.

We need to construct a relation over $P_1 \otimes Q$ and $P_2 \otimes Q$ and their derivative sets, and show that this is a system bisimulation. Let $P_i \stackrel{\text{def}}{=} \Sigma_i \bowtie_* \underline{\text{init}}.Con_i$ for $i = 1, 2$ and $Q \stackrel{\text{def}}{=} \Sigma \bowtie_* \underline{\text{init}}.Con$. Then by definition, $P_i \otimes Q \stackrel{\text{def}}{=} (\Sigma_i \bowtie_* \Sigma) \bowtie_* \underline{\text{init}}.(Con_i \bowtie_* Con)$ for $i = 1, 2$. By Proposition 7.1, these are well-defined and by Proposition 5.2, we know that

$$ds(P_i \otimes Q) = \{ \langle (\Sigma_i \bowtie_* \Sigma) \bowtie_* (Con_i \bowtie_* Con), \tau_i \rangle, \langle (\Sigma_i \bowtie_* \Sigma) \bowtie_* D_{i,1}, \sigma_{i,1} \rangle, \dots, \langle (\Sigma_i \bowtie_* \Sigma) \bowtie_* D_{i,n}, \sigma_{i,n} \rangle \}$$

where, for arbitrary state σ , $ds(\underline{\text{init}}.(Con_i \boxtimes Con)) = \{\langle Con_i \boxtimes Con, \sigma \rangle, \langle D_{i,1}, \sigma \rangle, \dots, \langle D_{i,n}, \sigma \rangle\}$. Moreover, since \boxtimes is a static operator, each $D_{i,k}$ has the form $C_{i,k} \boxtimes C_j$ for $C_{i,k} \in ds(\underline{\text{init}}.Con_i)$ and $C_j \in ds(\underline{\text{init}}.Con)$. We define the relation

$$\mathcal{R}' = \{((\Sigma_1 \boxtimes \Sigma) \boxtimes (C_{1,j_1} \boxtimes C_j), (\Sigma_2 \boxtimes \Sigma) \boxtimes (C_{2,j_2} \boxtimes C_j)) \mid \Sigma_1 \boxtimes C_{1,j_1} \sim_s \Sigma_2 \boxtimes C_{2,j_2}\}$$

and let $\mathcal{R} = \mathcal{R}' \cup \{(P_1 \otimes Q, P_2 \otimes Q)\}$. We will show that \mathcal{R} is a system bisimulation.

Consider a transition $\langle (\Sigma_1 \boxtimes \Sigma) \boxtimes (C_{1,j_1} \boxtimes C_j), \rho \rangle \xrightarrow{\underline{a}} \langle (\Sigma_1 \boxtimes \Sigma) \boxtimes (C_{1,k_1} \boxtimes C_k), \rho' \rangle$. By our assumption that all events appear in the controller, we can conclude that \underline{a} is a synchronised event in both the uncontrolled system and the controller. Thus we have 3 subcases: either the events come from Σ_1 and C_{1,j_1} , or from Σ and C_j , or from $\Sigma_1 \boxtimes \Sigma$ and $C_{1,j_1} \boxtimes C_j$. We will prove the last case, as the other two are simpler.

By shorter inferences in the derivation tree of the transition, we can conclude that $\langle \Sigma_1, \rho \rangle \xrightarrow{\underline{a}} \langle \Sigma_1, \mu \rangle$, $\langle \Sigma, \rho \rangle \xrightarrow{\underline{a}} \langle \Sigma, \nu \rangle$, $\langle C_{1,j_1}, \rho \rangle \xrightarrow{\underline{a}} \langle C_{1,k_1}, \rho \rangle$ and $\langle C_j, \rho \rangle \xrightarrow{\underline{a}} \langle C_k, \rho \rangle$, with $\rho' = \Gamma(\rho, \Gamma(\rho, \mu, \nu), \rho) = \Gamma(\rho, \mu, \nu)$ where ρ' is the state appearing in the target configuration of the original transition. Therefore $\langle \Sigma_1 \boxtimes C_{1,j_1}, \rho \rangle \xrightarrow{\underline{a}} \langle \Sigma_1 \boxtimes C_{1,k_1}, \mu \rangle$ since $\Gamma(\rho, \mu, \rho) = \mu$.

Since $\Sigma_1 \boxtimes C_{1,j_1} \sim_s \Sigma_2 \boxtimes C_{2,j_2}$, there is a transition $\langle \Sigma_1 \boxtimes C_{2,j_2}, \rho \rangle \xrightarrow{\underline{a}} \langle \Sigma_1 \boxtimes C_{2,k_2}, \mu \rangle$ such that $\Sigma_1 \boxtimes C_{1,k_1} \sim_s \Sigma_2 \boxtimes C_{2,k_2}$. By a shorter inference, $\langle \Sigma_1, \rho \rangle \xrightarrow{\underline{a}} \langle \Sigma_1, \mu \rangle$ and $\langle C_{2,j_2}, \rho \rangle \xrightarrow{\underline{a}} \langle C_{2,k_2}, \rho \rangle$. This allows us to infer the transition $\langle (\Sigma_2 \boxtimes \Sigma) \boxtimes (C_{2,j_2} \boxtimes C_j), \rho \rangle \xrightarrow{\underline{a}} \langle (\Sigma_2 \boxtimes \Sigma) \boxtimes (C_{2,k_2} \boxtimes C_k), \rho' \rangle$ as required. The symmetric condition is proved similarly. The same approach can be used to show that the pair $(P_1 \otimes Q, P_2 \otimes Q)$ are bisimilar by considering their init transitions (they have no others). \square

The following result shows that it is the prefixes which determine the behaviour of the controlled systems because of the restrictions imposed on well-defined controlled systems.

Theorem 8.3. Let $\Sigma_1 \boxtimes_L \underline{\text{init}}.Con$ and $\Sigma_2 \boxtimes_L \underline{\text{init}}.Con$ be two well-defined controlled systems.

If $\text{pr}(\Sigma_1) = \text{pr}(\Sigma_2)$ then $\Sigma_1 \boxtimes_L \underline{\text{init}}.Con \sim_s \Sigma_2 \boxtimes_L \underline{\text{init}}.Con$

Proof. We will show that $\{(\Sigma_1, \Sigma_2)\}$ is a bisimulation. Since we are dealing with well-defined controlled systems, we know that each prefix occurs in a subcomponent of the form $S \stackrel{\text{def}}{\underline{a}}: (t, r, I(\vec{X})).S + \dots$ and this means that \underline{a} can always happen. Additionally, for any other prefix starting with \underline{a} appearing in a different subcomponent, these events will be synchronised on. Blocking due to a subcomponent being unable to perform the same action cannot occur because all events are able to occur. Hence the events that Σ_1 can perform are all the events that appear in a prefix in $\text{pr}(\Sigma_1)$. Because of the form of S , we have that $\Sigma_1 \xrightarrow{\underline{a}} \Sigma_1$ for each possible event \underline{a} . This argument applies for each \underline{a} in a prefix in $\text{pr}(\Sigma_2)$. Since $\text{pr}(\Sigma_1) = \text{pr}(\Sigma_2)$, $\{(\Sigma_1, \Sigma_2)\}$ is a bisimulation and $\Sigma_1 \sim_s \Sigma_2$. Since \sim_s is a congruence, we have the result. \square

This result allows us to tell whether two HYPE models are bisimilar by inspecting the prefixes in the model description to see if they are the same. Hence bisimulation can be checked syntactically. The next two results show that bisimilar HYPE models have the same ODEs.

Lemma 8.1. Let P and Q be well-defined controlled systems. If $P \sim_s Q$ then $\text{st}(P) = \text{st}(Q)$.

Proof. Consider $\langle P', \sigma' \rangle$ a derivative of $\langle P, \sigma \rangle$ then $\sigma' \in \text{st}(P)$. Since $P \sim_s Q$, we can find $\langle Q', \sigma' \rangle$ with $\sigma' \in \text{st}(Q)$. Hence $\text{st}(P) \subseteq \text{st}(Q)$ and vice versa. \square

Proposition 8.1. Let P and Q be well-defined controlled systems. If $P \sim_s Q$ then for every state $\sigma \in \text{st}(P)$, $P_\sigma = Q_\sigma$.

Proof. The two well-defined controlled systems are $(P, \mathcal{V}, \mathcal{X}, \mathcal{E}, \mathcal{A}, ec, iv, pr, EC, ID)$ and $(Q, \mathcal{V}, \mathcal{X}, \mathcal{E}, \mathcal{A}, ec, iv, pr, EC, ID)$. Hence

$$P_\sigma = \left\{ \frac{dV}{dt} = \sum \{r[I(\vec{W})] \mid iv(t) = V \text{ and } \sigma(t) = (r, I(\vec{W}))\} \mid V \in \mathcal{V} \right\}$$

and

$$Q_\sigma = \left\{ \frac{dV}{dt} = \sum \{r[I(\vec{W})] \mid iv(t) = V \text{ and } \sigma(t) = (r, I(\vec{W}))\} \mid V \in \mathcal{V} \right\}$$

which are clearly the same. \square

The converse of Proposition 8.1 does not hold. It is possible for two states to be the same and hence give identical ODEs, but this does not mean that their associated derivatives are system bisimilar. We now go on to consider bisimulations that have been defined for other hybrid process algebras and those defined directly on hybrid systems.

8.1. Bisimulations for hybrid process algebras

Other process algebras for hybrid systems use a hybrid transition system with two types of transition: one type represents discrete events and the other continuous evolution of the system [Kha06]. By comparison, our transition system only has transitions for events. This gives a smaller transition system on which it is possible to consider simpler notions of equivalence.

Bergstra and Middelburg [BM05] present two bisimulations for their process algebra for hybrid systems ACP_{hs}^{srt} , defined over a hybrid transition system. One bisimulation fits with their axiomatic definition, and the other gives congruence with respect to the parallel operator, and was originally defined in [CR03]. We can recast these equivalences for our transition system and for our language, so that they are of interest here.

Definition 8.3 (Bergstra and Middelburg 2005). A *bisimulation* is a binary relation $B \subseteq (\mathcal{C} \times \mathcal{S}) \times (\mathcal{C} \times \mathcal{S})$ such that for all $(\langle P, \sigma \rangle, \langle Q, \sigma \rangle) \in B$, for all $\underline{a} \in \mathcal{E}$, whenever

1. $\langle P, \sigma \rangle \xrightarrow{\underline{a}} \langle P', \sigma' \rangle$, there exists $\langle Q', \sigma' \rangle$ such that $\langle Q, \sigma \rangle \xrightarrow{\underline{a}} \langle Q', \sigma' \rangle$ and $(\langle P', \sigma' \rangle, \langle Q', \sigma' \rangle) \in B$, and
2. $\langle Q, \sigma \rangle \xrightarrow{\underline{a}} \langle Q', \sigma' \rangle$, there exists $\langle P', \sigma' \rangle$ such that $\langle P, \sigma \rangle \xrightarrow{\underline{a}} \langle P', \sigma' \rangle$ and $(\langle P', \sigma' \rangle, \langle Q', \sigma' \rangle) \in B$.

$\langle P, \sigma \rangle$ and $\langle Q, \sigma \rangle$ are *bisimilar*, written $\langle P, \sigma \rangle \dot{\simeq} \langle Q, \sigma \rangle$ if they are contained in a bisimulation. P and Q are *bisimilar*, written $P \dot{\simeq} Q$ if $\langle P, \sigma \rangle \dot{\simeq} \langle Q, \sigma \rangle$ for all states σ .

Definition 8.4 (Bergstra and Middelburg 2005). A *interference-compatible (ic-)bisimulation* is a binary relation $B \subseteq \mathcal{C} \times \mathcal{C}$ such that for all $(P, Q) \in B$, for all $\underline{a} \in \mathcal{E}$, for all $\sigma \in \mathcal{S}$ whenever

1. $\langle P, \sigma \rangle \xrightarrow{\underline{a}} \langle P', \sigma' \rangle$, there exists $\langle Q', \sigma' \rangle$ such that $\langle Q, \sigma \rangle \xrightarrow{\underline{a}} \langle Q', \sigma' \rangle$ and $(P', Q') \in B$, and
2. $\langle Q, \sigma \rangle \xrightarrow{\underline{a}} \langle Q', \sigma' \rangle$, there exists $\langle P', \sigma' \rangle$ such that $\langle P, \sigma \rangle \xrightarrow{\underline{a}} \langle P', \sigma' \rangle$ and $(P', Q') \in B$.

P and Q are *ic-bisimilar*, written $P \dot{\simeq}^i Q$ if they are contained in a ic-bisimulation.

It can be seen that $\dot{\simeq}^i$ is the same as \sim_s since our system bisimulation definition also requires identical states. Both bisimulation and ic-bisimulation are congruences for the sequential operators of ACP_{hs}^{srt} ; however, bisimulation is not a congruence for the parallel and merge operators. On the other hand, certain axioms are not sound for ic-bisimulation, hence the need for a complex two-level derivation system. In contrast, for HYPE, these two bisimulations identify the same models as we now show.

Theorem 8.4. For well-defined controlled systems P and Q , $P \dot{\simeq} Q$ implies $P \sim_s Q$.

Proof. Let $B = \{(P, Q) \mid P \dot{\simeq} Q\}$. Consider $(P, Q) \in B$ with $\langle P, \sigma \rangle \xrightarrow{\underline{a}} \langle P', \sigma' \rangle$. Since $P \dot{\simeq} Q$, there exists $\langle Q', \sigma' \rangle$ such that $\langle Q, \sigma \rangle \xrightarrow{\underline{a}} \langle Q', \sigma' \rangle$ and $\langle P', \sigma' \rangle \dot{\simeq} \langle Q', \sigma' \rangle$. If $P' \dot{\simeq} Q'$ then the proof is complete, so let $B' = \{(\langle P, \tau \rangle, \langle Q, \tau \rangle) \mid \langle P, \sigma \rangle \dot{\simeq} \langle Q, \sigma \rangle \text{ for some } \sigma\}$. Consider arbitrary τ' with $\langle P', \tau' \rangle \xrightarrow{\underline{a}} \langle P'', \tau'' \rangle$.

We can write $\tau'' = \tau'[\iota_1 \mapsto (r_1, I_1)] \dots [\iota_n \mapsto (r_n, I_n)]$ for appropriate n and $\iota_j, r_j, I_j, 1 \leq j \leq n$. Let $\sigma'' = \sigma'[\iota_1 \mapsto (r_1, I_1)] \dots [\iota_n \mapsto (r_n, I_n)]$. By Proposition 5.3, $\langle P', \sigma' \rangle \xrightarrow{\underline{a}} \langle P'', \sigma'' \rangle$ and since $\langle P', \sigma' \rangle \dot{\simeq} \langle Q', \sigma' \rangle$, there exists $\langle Q'', \sigma'' \rangle$ such that $\langle Q', \sigma' \rangle \xrightarrow{\underline{a}} \langle Q'', \sigma'' \rangle$ and $\langle P'', \sigma'' \rangle \dot{\simeq} \langle Q'', \sigma'' \rangle$. Moreover $\langle Q', \tau' \rangle \xrightarrow{\underline{a}} \langle Q'', \tau'' \rangle$ by Proposition 5.3 with $(\langle P'', \tau'' \rangle, \langle Q'', \tau'' \rangle) \in B'$. Hence $P' \dot{\simeq} Q'$ and $(P', Q') \in B$. The symmetric case is proved similarly. \square

Theorem 8.5. For well-defined controlled systems P and Q , $P \dot{\simeq}^i Q$ implies $P \dot{\simeq} Q$.

Proof. Let $B = \{(\langle P, \sigma \rangle, \langle Q, \sigma \rangle) \mid P \dot{\simeq}^i Q\}$. Consider $(\langle P, \sigma \rangle, \langle Q, \sigma \rangle) \in B$ with $\langle P, \sigma \rangle \xrightarrow{\underline{a}} \langle P', \sigma' \rangle$. Since $P \dot{\simeq}^i Q$, there exists $\langle Q', \sigma' \rangle$ such that $\langle Q, \sigma \rangle \xrightarrow{\underline{a}} \langle Q', \sigma' \rangle$ and $P' \dot{\simeq}^i Q'$, hence $(\langle P', \sigma' \rangle, \langle Q', \sigma' \rangle) \in B$ and B is a bisimulation. Since σ is arbitrary, $P \dot{\simeq} Q$. \square

The counter-example presented by Bergstra and Middelburg [BM05] to illustrate that congruence of bisimulation does not hold for a parallel operator uses the fact that in bisimulation, it is only required that target configurations are bisimilar for a specific state. This can be described as an initially stateless

bisimilarity [Kha08, MRG05] since it is only for the initial pair that they must be able to match all transitions for any state. Hence it is possible to have two bisimilar processes, but when each is put in parallel with another process, the state in the target configuration can change, allowing previously blocked actions to become available. In contrast, under ic-bisimulation, target configurations are only ic-bisimilar if they have the same transitions under all possible states. This is described as stateless bisimilarity, because all pairs must be able to match transitions for any state. One can also consider a state-based bisimulation where pairs must be able to match transitions for specific states only but this type of bisimulation is not considered here.

More specifically, one of the two processes involves the setting of the variable v to 1, followed by an a event prefixing deadlock, and the other has the same variable setting followed by event a prefixing a condition $v = 0$ that must be true for the event b to occur. These are equivalent under bisimulation since for any derived state, after the a event, the state contains the fact that v is associated with the value 1 and the b -event can never occur. However, when putting these two in parallel with a process that can change the value of v to zero, it is possible for the b event to occur in the second process but not the first. In the case of ic-bisimulation the two target states are not bisimilar since all possible states must be considered, including those where $v = 0$.

This cannot occur in HYPE because activation conditions (which determine if something can happen because a variable has the appropriate value) and resets (which change the values of variables) are disjoint from the syntactic description of the model, and the labelled transition system generated by the operational semantics does not take these into account. Thus, although bisimulation may only consider some states when applied to HYPE models, it is not possible for the same situation to arise.

8.2. Bisimulations on hybrid automata

The standard notion of bisimulation of hybrid automata is defined for a transition system encoding the dynamical evolution. Both continuous transitions and discrete transitions are used. These two relations are combined (usually interleaving discrete and continuous transitions) into one relation which defines the behaviour of the system. The hybrid automata bisimulation is defined on this relation in the usual way [Hen96, HTP05, DT07]. Most of the research into these bisimulations focusses on finding restrictions on hybrid automata syntax implying the existence of a finite bisimulation quotient thus guaranteeing decidability of reachability and of model checking, such as [LPS00].

Another form of bisimulation, the U -bisimulation, acts on the control graph, and it is used for projecting away some variables and for collapsing some modes of the automaton [AMP⁺03]. In the following, we consider a simplified form of this U -bisimulation, where we retain all variables and collapse modes that have the same structure.

Definition 8.5. Let $\mathcal{H} = (V, E, \mathbf{X}, \mathcal{E}, flow, init, inv, event, jump, reset, urgent)$ be an hybrid automaton. Two modes v_1, v_2 are U -bisimilar, $v_1 \sim_{HA} v_2$, if and only if the following conditions hold:

1. $flow(v_1) = flow(v_2)$, $inv(v_1) = inv(v_2)$, and $init(v_1) = init(v_2)$;
2. for each edge $e_1 = (v_1, v'_1)$ there exists an edge $e_2 = (v_2, v'_2)$ such that $event(e_1) = event(e_2)$, $reset(e_1) = reset(e_2)$, $jump(e_1) = jump(e_2)$, $urgent(e_1) = urgent(e_2)$, and $v'_1 \sim_{HA} v'_2$;
3. for each edge $e_2 = (v_2, v'_2)$ there exists an edge $e_1 = (v_1, v'_1)$ such that $event(e_1) = event(e_2)$, $reset(e_1) = reset(e_2)$, $jump(e_1) = jump(e_2)$, $urgent(e_1) = urgent(e_2)$, and $v'_1 \sim_{HA} v'_2$;

This notion of bisimulation is related to system bisimulation. We have the following lemma.

Proposition 8.2. Let $P_1, P_2 \in \mathcal{C}_{Sys}$ such that $P_1 \sim_s P_2$. Then $\langle P_1, \sigma \rangle \sim_{HA} \langle P_2, \sigma \rangle$.

Proof. The first point of Definition 8.5 follows because we are considering the same state σ , assuming the initial conditions to be the same (they also depend just on σ). The next two points follow from the definition of \sim_s and from the fact that the predicates $jump, event, urgent, reset$ depend just on the event labelling the edge. \square

Interestingly, the converse of the lemma does not hold. In fact, two states $\langle P, \sigma \rangle$ and $\langle Q, \tau \rangle$ of an hybrid automaton can be U -bisimilar even if $\sigma \neq \tau$: we only require the differential equations obtained from σ and τ to be the same. In fact, different states can lead to the same set of equations: consider two activities α and β mapped by iv to the same variable V . Then, if σ contains $(\alpha, 1, I)$ and $(\beta, -1, I)$ and τ contains $(\alpha, 2, I)$ and $(\beta, -2, I)$, then both σ and τ imply the equation $\dot{V} = 0$.

This means that the notion of bisimulation for hybrid automata is coarser than that for HYPE. Indeed, in a HYPE model we make explicit the source of each single flow of the system, while in a hybrid automaton flows are merged together in differential equations, and they cannot be separated in single influences. Stated otherwise, ODEs irremediably lose information about the logic of the system.

9. Modelling Style

In this section we discuss in more detail the modelling style of HYPE, comparing it with the modelling style of other hybrid process algebras and hybrid automata.

The main difference between HYPE and other hybrid process algebras is in the description of the continuous dynamics. Most of these process algebras describe the continuous dynamics by explicitly providing ODEs for each variable [BM05, CR05, RS03, vBMR⁺06] and this is also the case for hybrid automata [Hen96]. In HYPE, we have adopted a different approach, modelling a set of influences, or flows, that affect the evolution of continuous variables, and then deriving the ODE by looking at the set of active influences in each mode. Discrete modes themselves are not modelled explicitly, but are derived from the operational semantics of the language. Since this represents a significant change in style from the existing process algebras, it was not feasible to develop HYPE as an extension of an existing hybrid process algebra. Instead we decided that HYPE must be developed afresh, centred on the idea of additive influences.

The advantage of this mechanism to describe continuous dynamics is particularly evident for systems in which there are many influences that can affect the dynamics of a specific variable, each of them having two or more possible forms. Consider again the orbiter example of Section 2 (see also Figure 1), and focus on the continuous dynamics of the temperature K . This variable is affected by four different flow sources: the heat of the sun, the heater, the effect of the shade and the thermodynamic cooling. The first three sources can be in two different states (i.e. present or absent), and their effect on the temperature depends on their state. Indeed, this gives rise to eight different differential equations for the temperature K . In languages that require the explicit provision of the ODEs for each variable, and that do not have a compositional mechanisms at the level of the continuous dynamics, like hybrid automata [Hen96] and ACP_{hs}^{srt} [BM05], the modeller is forced to write these eight ODEs explicitly. We can see this in Figure 13, where an ACP_{hs}^{srt} model for the orbiter is presented. This monolithic representation of the ODEs would be similar in any of the other existing hybrid process algebras. Notice that the number of possible ODEs grows exponentially with the number of flow sources that have more than one state (think of a tank with many inflow and outflow pipes). In HYPE, instead, we model the continuous dynamics of the orbiter temperature by associating an influence with each source, and changing the form of the influence in response to discrete events (i.e. activation and inactivation of the flow sources). This allows a much more compact description of the continuous dynamics: the complexity of the description grows linearly with the number of flow sources.

The use of additive flows also has the advantage that HYPE has more expressive power in relation to compositionality, as demonstrated by Theorem 7.1 and Corollary 7.1. In these results, it is shown that hybrid automata product is restricted to the case where there are no shared variables between the two hybrid automata (or flows are zero for all variables), whereas the product of HYPE models allows for shared variables. This means that it is possible to add new influences to an existing model that affect one or more of that model's variables. To achieve a similar effect in a hybrid automaton, the flow at some modes must be changed directly since it cannot be achieved using a product construction. In HYPE, addition of new levels can either be done at the model level using product, or as we demonstrate in the next section, by adding new subcomponents with new influences.

HYPE is also characterised by modularity in its specification style. Subcomponents are defined separately from controllers and the two have events on which they synchronise. Additionally, event conditions are specified separately from events, influence names from variables, and influence strength from influence type. In general, this means that a modeller can consider different concerns separately and we now highlight two examples of this modularity, and show the practical benefits it can bring. The reader can observe how the different states of the influences for the orbiter (see Figure 1) differ only in their influence strength, and not in their influence type. We found this to be a common pattern, and this motivated our separation of the influence strength from the influence type, i.e. the description of the functional form of the flow from its intensity. Moreover, as discussed below, we chose to combine influences into ODEs in an additive way.

Another modular feature of HYPE, is the strong separation that exists between the description of the continuous dynamics and the description of the discrete control. In particular, in HYPE one models the

$$\begin{aligned}
T_{\mathbf{dFD}} &\stackrel{\text{def}}{=} (\dot{K} = -K \wedge K \geq t_2 \wedge K \leq t_3) \triangleright \sigma_{\text{rel}}^* \left(((K^\bullet = \bullet K) \triangleright \widetilde{\text{light}} \cdot T_{\mathbf{1FD}}) + \right. \\
&\quad \left. ((K \leq t_2) \rightarrow ((K^\bullet = \bullet K) \triangleright \widetilde{\text{on}} \cdot T_{\mathbf{dND}})) + ((K \geq t_3) \rightarrow ((K^\bullet = \bullet K) \triangleright \widetilde{\text{up}} \cdot T_{\mathbf{dNU}})) \right) \\
T_{\mathbf{1FD}} &\stackrel{\text{def}}{=} (\dot{K} = r_s - K \wedge K \geq t_2 \wedge K \leq t_3) \triangleright \sigma_{\text{rel}}^* \left(((K^\bullet = \bullet K) \triangleright \widetilde{\text{dark}} \cdot T_{\mathbf{dFD}}) + \right. \\
&\quad \left. ((K \leq t_2) \rightarrow ((K^\bullet = \bullet K) \triangleright \widetilde{\text{on}} \cdot T_{\mathbf{dND}})) + ((K \geq t_3) \rightarrow ((K^\bullet = \bullet K) \triangleright \widetilde{\text{up}} \cdot T_{\mathbf{dNU}})) \right) \\
T_{\mathbf{dND}} &\stackrel{\text{def}}{=} (\dot{K} = r_h - K \wedge K \leq t_1 \wedge K \leq t_3) \triangleright \sigma_{\text{rel}}^* \left(((K^\bullet = \bullet K) \triangleright \widetilde{\text{light}} \cdot T_{\mathbf{1ND}}) + \right. \\
&\quad \left. ((K \geq t_1) \rightarrow ((K^\bullet = \bullet K) \triangleright \widetilde{\text{off}} \cdot T_{\mathbf{dFD}})) + ((K \geq t_3) \rightarrow ((K^\bullet = \bullet K) \triangleright \widetilde{\text{up}} \cdot T_{\mathbf{dNU}})) \right) \\
T_{\mathbf{1ND}} &\stackrel{\text{def}}{=} (\dot{K} = r_s + r_h - K \wedge K \leq t_1 \wedge K \leq t_3) \triangleright \sigma_{\text{rel}}^* \left(((K^\bullet = \bullet K) \triangleright \widetilde{\text{dark}} \cdot T_{\mathbf{dND}}) + \right. \\
&\quad \left. ((K \geq t_1) \rightarrow ((K^\bullet = \bullet K) \triangleright \widetilde{\text{off}} \cdot T_{\mathbf{1FD}})) + ((K \geq t_3) \rightarrow ((K^\bullet = \bullet K) \triangleright \widetilde{\text{up}} \cdot T_{\mathbf{1NU}})) \right) \\
T_{\mathbf{dFU}} &\stackrel{\text{def}}{=} (\dot{K} = -r_d - K \wedge K \geq t_2 \wedge K \geq t_4) \triangleright \sigma_{\text{rel}}^* \left(((K^\bullet = \bullet K) \triangleright \widetilde{\text{light}} \cdot T_{\mathbf{1FU}}) + \right. \\
&\quad \left. ((K \leq t_2) \rightarrow ((K^\bullet = \bullet K) \triangleright \widetilde{\text{on}} \cdot T_{\mathbf{dNU}})) + ((K \leq t_4) \rightarrow ((K^\bullet = \bullet K) \triangleright \widetilde{\text{down}} \cdot T_{\mathbf{dND}})) \right) \\
T_{\mathbf{1FU}} &\stackrel{\text{def}}{=} (\dot{K} = r_s - r_d - K \wedge K \geq t_2 \wedge K \geq t_4) \triangleright \sigma_{\text{rel}}^* \left(((K^\bullet = \bullet K) \triangleright \widetilde{\text{dark}} \cdot T_{\mathbf{dFU}}) + \right. \\
&\quad \left. ((K \leq t_2) \rightarrow ((K^\bullet = \bullet K) \triangleright \widetilde{\text{on}} \cdot T_{\mathbf{dNU}})) + ((K \leq t_4) \rightarrow ((K^\bullet = \bullet K) \triangleright \widetilde{\text{down}} \cdot T_{\mathbf{dND}})) \right) \\
T_{\mathbf{dNU}} &\stackrel{\text{def}}{=} (\dot{K} = r_h - r_d - K \wedge K \leq t_1 \wedge K \geq t_4) \triangleright \sigma_{\text{rel}}^* \left(((K^\bullet = \bullet K) \triangleright \widetilde{\text{light}} \cdot T_{\mathbf{1NU}}) + \right. \\
&\quad \left. ((K \geq t_1) \rightarrow ((K^\bullet = \bullet K) \triangleright \widetilde{\text{off}} \cdot T_{\mathbf{dFU}})) + ((K \leq t_4) \rightarrow ((K^\bullet = \bullet K) \triangleright \widetilde{\text{down}} \cdot T_{\mathbf{dND}})) \right) \\
T_{\mathbf{1NU}} &\stackrel{\text{def}}{=} (\dot{K} = r_s + r_h - r_d - K \wedge K \leq t_1 \wedge K \geq t_4) \triangleright \sigma_{\text{rel}}^* \left(((K^\bullet = \bullet K) \triangleright \widetilde{\text{dark}} \cdot T_{\mathbf{dNU}}) + \right. \\
&\quad \left. ((K \geq t_1) \rightarrow ((K^\bullet = \bullet K) \triangleright \widetilde{\text{off}} \cdot T_{\mathbf{1FU}})) + ((K \leq t_4) \rightarrow ((K^\bullet = \bullet K) \triangleright \widetilde{\text{down}} \cdot T_{\mathbf{1ND}})) \right)
\end{aligned}$$

Fig. 13. The heater component expressed in ACP_{hs}^{srt} where the first element of the subscript indicates light or dark, the second of **F** or **oN**, and the third **Down** or **Up**.

uncontrolled system and the controller, as distinct processes. The uncontrolled system is basically a flat description of all possible states of all flow sources, while the controller is used to impose an ordering on discrete events. The dependence of discrete events on continuous variables, via activation conditions and resets, is described separately from the controller. This allows the modeller to experiment with different combinations of discrete controllers, by either changing causality dependencies between events or by changing their dependencies on the continuous part of the system. Moreover, the choice of describing event conditions separately from the controlled system allows us to separate modelling concerns. In particular, we separate the description of the architecture of the model, i.e. the logical structure defining the possible interactions, from the description of the quantitative features of dynamics. This allows the modeller to experiment with different dynamical regimes without the necessity of modifying the structure of the model. We stress the fact that by combining the controller with activation conditions and resets, one can describe arbitrarily complex discrete dynamics.

Our choice to combine influences in an additive way when forming the ODEs, stems from the identification of influences with flow sources, which tend to be additive in nature (think of the inflows and outflows of a tank, or of the combined effect of different chemical reactions in a chemical system). We wish to stress, however, that this additive style is not a limitation. In fact, we can obtain any set of ODEs, with nonlinear terms as complex as desired. Indeed, this can be achieved by appropriate definition of the influence types. In the extreme case, one can model an ODE $\dot{V} = f(V)$ for a variable V by a single influence with strength 1 and influence type $I(V) = f(V)$. However, whilst this is within the definition of the language, it is not inkeeping with the modelling spirit of HYPE, which emphasises modularity. It remains to remark that, in all the examples we have considered so far, additivity of flows has not been a limitation but we remain open to the possibility of studying different mechanisms to combine flows in the future.

In Section 4, we showed how to associate a hybrid automaton with each HYPE model. It is interesting to further consider the relationship between the class of models that can be described in HYPE and the class of models that can be described by hybrid automata (at least the class of hybrid automata that we

consider in in Definition 4.1). It is also possible to construct a HYPE model that is equivalent to any given hybrid automaton⁷. The only element of hybrid automata that is not straightforwardly described in HYPE is the non-urgent discrete transition (with non-trivial activation conditions). However, these transitions can be modelled in HYPE by a sequence of events: an urgent event to fire as soon as the guard of the transition becomes active, followed by a non-deterministic event that is disallowed when the guard is false. We do not provide formal details of the construction of a HYPE model from a hybrid automata here. An example of it is given by the train-gate controller model of Section 10.

A modelling approach that is closely related to HYPE in the way of dealing with continuous dynamics is the (hybrid) bond graph approach [CBM08]. However, bond graphs are tailored to model physical systems, while HYPE has a more abstract approach, that supports the modelling of a large number of classes of systems showing hybrid dynamics (e.g. biochemical systems, ecological systems, financial markets, social networks, computer networks, and so on), even if it can be less efficient in specific domains than languages designed for these domains.

We close this section commenting on the notion of state in HYPE and in other hybrid languages. Usually, in hybrid languages a state describes the current mode and the current evaluation of continuous variables. In HYPE, instead, a state (operational state) is just a collection of active influences, hence it models the shape of the continuous dynamics, rather than a specific configuration of the system. This difference reflects on the notions of bisimulation. While in hybrid languages like ACP_{hs}^{srt} or hybrid automata, bisimulations predicate properties of the systems dynamics, in HYPE bisimulation compares two models in terms of their structure. In this sense, it is a stronger notion of bisimulation, as models with non-bisimilar structure can still exhibit a bisimilar dynamical behaviour. HYPE bisimulation, in particular, can be used to reason about the complexity of the description of a model and about the causal relationships between events, while bisimulation on the dynamics can be used to carry out analysis at that level, for instance reachability analysis. Further illustrating HYPE's more abstract style of modelling, as demonstrated in Section 8, is the fact that for HYPE, stateless bisimilarity is the same as initially stateless bisimilarity. This difference reflects also on the intrinsic computational complexity of these two notions: while bisimulation for finite HYPE models is (efficiently) computable, as it can be reduced to a bisimulation for labelled graphs [DPP04], bisimulations at the dynamical level are usually undecidable [HKPV95].

10. Example: train gate controller

We now consider the train gate system as described in [AHH96, BM05, Kha08]. We consider a section of railway track with a crossing gate. The track starts 1500m or more before the gate and ends 100m afterwards. There is a sensor at 1000m before the gate that sends a signal to the controller that a train is approaching, and a sensor at 100m after the gate that sends an *exit* signal to the controller. The train travels at between 48m/s and 52m/s and after the first sensor is passed, between 40m/s and 52m/s. The gate, on receipt of a *lower* signal, changes its angle to the ground from 90° to 0° at a rate of -20°/s until it is closed. Likewise, on the receipt of a *raise* signal, changes its angle to the ground from 0° to 90° at a rate of 20°/s until it is open. It must always respond to *lower* and *raise* signals. The controller on the receipt of an *approach* signal, sends a *lower* signal to the gate within 5 seconds, and on the receipt of an *exit* signal, sends a *raise* signal to the gate within 5 seconds, assuming it does not receive an *approach* signal within that time. It must always be available to receive an *approach* or *exit* signal. Once a train clears the section of track (passes the sensor after the gate), a subsequent train may arrive at the start of the section of track.

Initially, the train is at 1400m before the gate, and the gate is open. The safety property of the system in which we are interested, is whether the gate is down when the train is within a certain distance of the gate, say 100m. We will also consider a modification of the model, where the train is able to travel faster, with a reduction of speed closer to the gate, and consider whether the safety property still holds.

The HYPE model is given in Figure 14. Since HYPE does not directly support ranges of resets or ranges of influences, we will make decisions to simplify our modelling task. We will assume that subsequent trains appear at 1500m before the gate, that trains travel at their maximum speed of 52m/s and that the controller

⁷ The notion of equivalence we have in mind here is the following: given a hybrid automaton, we associate with it a HYPE model. Then, we construct a new hybrid automaton associated with this HYPE model according to the construction of Section 4, and prove that it is bisimilar to the original one.

$$\begin{aligned}
 \text{Gate} &\stackrel{\text{def}}{=} \underline{\text{lower}} : (g, -r_{gt}, c). \text{Gate} + \underline{\text{closed}} : (g, 0, c). \text{Gate} + \\
 &\quad \underline{\text{raise}} : (g, -r_{gt}, c). \text{Gate} + \underline{\text{open}} : (g, 0, c). \text{Gate} + \\
 &\quad \underline{\text{init}} : (g, 0, c). \text{Gate} \\
 \\
 \text{Train} &\stackrel{\text{def}}{=} \underline{\text{init}} : (d, r_{tr}, c). \text{Train} \\
 \\
 \text{Timer}_L &\stackrel{\text{def}}{=} \underline{\text{appr}} : (t_L, 1, \text{const}). \text{Timer}_L + \underline{\text{lower}} : (t_L, 0, \text{const}). \text{Timer}_L + \underline{\text{init}} : (t_L, 0, \text{const}). \text{Timer}_L \\
 \text{Timer}_R &\stackrel{\text{def}}{=} \underline{\text{exit}} : (t_R, 1, \text{const}). \text{Timer}_R + \underline{\text{raise}} : (t_R, 0, \text{const}). \text{Timer}_R + \underline{\text{init}} : (t_R, 0, \text{const}). \text{Timer}_R \\
 \\
 \text{Con}_a &\stackrel{\text{def}}{=} \underline{\text{appr}}. \text{Con}_l + \underline{\text{exit}}. \text{Con}_a & \text{Con}_l &\stackrel{\text{def}}{=} \underline{\text{appr}}. \text{Con}_l + \underline{\text{exit}}. \text{Con}_l + \underline{\text{lower}}. \text{Con}_e \\
 \text{Con}_e &\stackrel{\text{def}}{=} \underline{\text{appr}}. \text{Con}_e + \underline{\text{exit}}. \text{Con}_{ra} & \text{Con}_{ra} &\stackrel{\text{def}}{=} \underline{\text{appr}}. \text{Con}_l + \underline{\text{exit}}. \text{Con}_{ra} + \underline{\text{raise}}. \text{Con}_a \\
 \\
 \text{GC}_o &\stackrel{\text{def}}{=} \underline{\text{raise}}. \text{GC}_o + \underline{\text{lower}}. \text{GC}_l & \text{GC}_l &\stackrel{\text{def}}{=} \underline{\text{raise}}. \text{GC}_r + \underline{\text{lower}}. \text{GC}_l + \underline{\text{closed}}. \text{GC}_c \\
 \text{GC}_c &\stackrel{\text{def}}{=} \underline{\text{raise}}. \text{GC}_r + \underline{\text{lower}}. \text{GC}_c & \text{GC}_r &\stackrel{\text{def}}{=} \underline{\text{raise}}. \text{GC}_r + \underline{\text{lower}}. \text{GC}_l + \underline{\text{open}}. \text{GC}_o \\
 \\
 \text{Seq}_a &\stackrel{\text{def}}{=} \underline{\text{appr}}. \text{Seq}_p & \text{Seq}_p &\stackrel{\text{def}}{=} \underline{\text{pass}}. \text{Seq}_e & \text{Seq}_e &\stackrel{\text{def}}{=} \underline{\text{exit}}. \text{Seq}_a \\
 \\
 L &= \{\underline{\text{lower}}, \underline{\text{raise}}, \underline{\text{appr}}, \underline{\text{pass}}, \underline{\text{exit}}, \underline{\text{init}}\} \\
 \text{System} &\stackrel{\text{def}}{=} (\text{Gate} \boxtimes_{\{\underline{\text{init}}\}} \text{Train} \boxtimes_{\{\underline{\text{init}}\}} \text{Timer}_L \boxtimes_{\{\underline{\text{init}}\}} \text{Timer}_R) \boxtimes_L \underline{\text{init}}. ((\text{Con}_a \boxtimes_{\{\underline{\text{appr}}, \underline{\text{exit}}\}} \text{Seq}_a) \boxtimes_{\{\underline{\text{raise}}, \underline{\text{lower}}\}} \text{GC}_o) \\
 \\
 iv(g) &= G & iv(d) &= D & iv(t_L) &= T_L & iv(t_R) &= T_R & \llbracket c \rrbracket &= 1 \\
 \\
 ec(\underline{\text{init}}) &= (\text{true}, G' = 90 \wedge D' = -1400 \wedge T'_L = 0 \wedge T'_R = 0) \\
 \\
 ec(\underline{\text{lower}}) &= (T_L = 5, T'_L = 0) & ec(\underline{\text{closed}}) &= (G = 0, \text{true}) \\
 ec(\underline{\text{raise}}) &= (T_R = 5, T'_R = 0) & ec(\underline{\text{open}}) &= (G = 90, \text{true}) \\
 \\
 ec(\underline{\text{appr}}) &= (D = -1000, T'_L = 0) & r_{gt} &= 20 \\
 ec(\underline{\text{pass}}) &= (D = 0, \text{true}) & r_{tr} &= 52 \\
 ec(\underline{\text{exit}}) &= (D = 100, T'_R = 0 \wedge D' = -1500)
 \end{aligned}$$

Fig. 14. HYPE model of train gate controller

will take the full 5 seconds to signal the gate, since these values represent the most negative conditions for the system. After our model description, we will demonstrate how to modify it to allow more flexibility.

In the model, there are 4 continuous components to represent the train, the gate, and two timers. There are also 3 distinct controlling or sequencing elements. Seq_a captures the fact that the train does not ever go backwards. Con_a is the controller of the system, and GC_o represents the discrete aspects of the gate itself. The gate is required to accept all raise and lower signals but it is reasonable to expect that the gate has a notion of internal state to determine what to do if a signal is received, and is aware when it is closed or open so that it can stop its own motors. All of these aspects given here are also present in the hybrid automata model given in [AHH96] but in our model continuous and discrete aspects are separated.

The model is well-defined and we wish to check that it is also well-behaved. Since the events of GC_o which do not appear in $\text{Con}_a \boxtimes_* \text{Seq}_a$ do not activate any of those in $\text{Con}_a \boxtimes_* \text{Seq}_a$ and *vice versa*, we can apply Proposition 6.4. It is necessary to consider $\text{Con}_a \boxtimes_* \text{Seq}_a$ as a single system since Con_a does allow infinite instantaneous behaviour. The event appr does not inhibit itself, and there are loops involving this event in the definition of Con_a . In the I-graph of Con_a , we obtain the loop $(\text{Con}_l, \underline{\text{appr}}, 10001) \rightarrow (\text{Con}_l, \underline{\text{appr}}, 10001)$ where the vector order is appr, pass, exit, lower, raise.

The I-graph of GC_o consists of four directed acyclic graphs (when only considering those nodes that are reachable from nodes with vectors consisting of all ones), and the I-graph of $\text{Con}_a \boxtimes_* \text{Seq}_a$ consists of 12 directed acyclic graphs, 6 of which are single nodes and 6 are two node graphs. Since these are both acyclic, we know that the overall controller is well-behaved by Proposition 6.4. There are 22 configurations

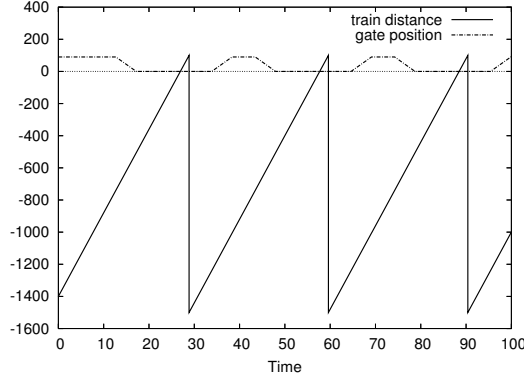


Fig. 15. Trace for the train gate system where $r_{tr} = 52$.

(excluding the initial configuration) in the labelled transition system of the model, hence giving a hybrid automaton with that number of modes.

When comparing our model to that of the hybrid automaton in [AHH96], our controller has an additional state reflecting a slightly more complex controller. If we construct the three-state controller, say Con' and put each controller in parallel with Seq_a , then we can show that $Con' \bowtie_* Seq_a$ and $Con_a \bowtie_* Seq_a$ are system bisimilar. From this we know by congruence, that $(Con' \bowtie_* Seq_a) \bowtie_* GC_o$ and $(Con_a \bowtie_* Seq_a) \bowtie_* GC_o$ are system bisimilar. Finally by Corollary 8.1, the two controlled systems are system bisimilar.

In terms of the safety property, we can run a single trace of the model (as given in Figure 15) since it is deterministic, and determine that the gate closes when the train is at 505.9m before the crossing agreeing with previous calculations [Kha08]. Another approach to considering the safety property is to construct a HYPE model with the same continuous components and a different controller that exactly describes how we want the system to behave with an additional event for failure. We address this below after we consider how to change the model.

To show the compositionality of our approach, we consider the scenario where the train is now able to travel at a much faster speed but the crossing gate is unchanged. To ensure that the gate is closed by the time the train is 100m from the crossing, it is necessary to slow the train as it approaches the crossing. We assume the train is required to slow at 1000m before the crossing.

Since we already have an event at $D = -1000$, we do not need to introduce another event. We can simply add a new train component,

$$TrainRS \stackrel{def}{=} \text{appr} : (d_s, -r_{sl}, c). TrainRS + \text{exit} : (d_s, 0, c). TrainRS + \text{init} : (d_s, 0, c). TrainRS$$

and let $iv(d_s) = D$. By adding this component, the train speed will be slowed at 1000m to $r_{tr} - r_{sl}$. In the model with the added component, we still have the same number of states in the labelled transition system because no new events have been added. The underlying hybrid automaton will be very similar to the previous one, except in certain states, the ODE for distance will differ slightly. In Figure 16, the graph for $r_{gt} = 140$ and $r_{sl} = 20$ is given on the left. We can see from the graph that the gate is not closed when the train passes the crossing and hence it cannot be closed 100m before the crossing. We wish to add to our model to capture the fact that this behaviour is not desirable.

We do this by adding an event called `fail` with event conditions $ec(\text{fail}) = (-100 \leq D \wedge G \neq 0, true)$. Since this event can happen at anytime, we add the controller, $FC \stackrel{def}{=} \text{fail}.0$. In the labelled transition system, this adds an additional configuration for each existing configuration (excluding the initial configuration). For each configuration $\langle \Sigma \bowtie_* D, \sigma \rangle$, there is a new configuration $\langle \Sigma \bowtie_* 0, \sigma \rangle$ with a new transition `fail` from the existing configuration to the new one. Thus the new hybrid automaton has 44 modes. This doubling of configurations and modes can be avoided by adding a `fail` event to each subcomponent and setting each influence to zero on the occurrence of the event. This means that there will be only one new state and there will be `fail` transitions from all existing states to the new state. The right graph in Figure 16 shows a trace of the gate system under the new speeds with the `fail` event being triggered.

By experimentation, we find that $r_{sl}=50$ leads to an execution without failure as shown in the left graph in Figure 17, however, this means that the gate is continually closed making the crossing unusable. By




Fig. 16. Traces for the train gate system where $r_{tr} = 140$ and $r_{sl} = 20$. The model in the left trace does not deal explicitly with failure whereas the model in the right trace has an explicit failure event




Fig. 17. Trace for the train gate system where $r_{tr} = 140$ and $r_{sl} = 50$ (left). Trace for the train gate system where $r_{tr} = 140$ and $r_{sl} = 50$, and originating distance is increased to 2000m

changing the gap between trains by increasing the originating distance to -2000m, we obtain a scenario where the behaviour of the gate is correct and the gate is opening between trains, as shown in the right graph in Figure 17.

Another approach to assessing the correctness of the model is to retain the continuous components of the system and use a controller that only allows the correct behaviour together with a failure component as before. We give the model in Figure 18. Using the underlying hybrid automaton which is smaller with only 10 modes, we obtain the same graphs as the last three using the same parameters.

We now consider aspects of the original scenario which we were not able to capture in HYPE. First, with

$$\begin{array}{ll}
 C_1 \stackrel{def}{=} \text{appr}.C_2 & C_6 \stackrel{def}{=} \text{appr}.C_7 + \text{raise}.C_8 \\
 C_2 \stackrel{def}{=} \text{lower}.C_3 & C_7 \stackrel{def}{=} \text{lower}.C_4 \\
 C_3 \stackrel{def}{=} \text{close}.C_4 & C_8 \stackrel{def}{=} \text{open}.C_1 + \text{appr}.C_9 \\
 C_4 \stackrel{def}{=} \text{pass}.C_5 & C_9 \stackrel{def}{=} \text{open}.C_2 + \text{lower}.C_3 \\
 C_5 \stackrel{def}{=} \text{exit}.C_6 & \\
 \\
 L = \{\text{lower}, \text{raise}, \text{appr}, \text{pass}, \text{exit}, \text{init}\} & \\
 \text{System} \stackrel{def}{=} (\text{Gate} \underset{\{\text{init}\}}{\bowtie} \text{Train} \underset{\{\text{init}\}}{\bowtie} \text{Timer}_L \underset{\{\text{init}\}}{\bowtie} \text{Timer}_R) \underset{L}{\bowtie} \text{init}.(C_1 \parallel FC) &
 \end{array}$$

Fig. 18. HYPE model of train gate controller with fail mode

respect to the originating distance being at least 1500m, we can use the \perp activation condition in the event condition for `exit`. This simulates a train starting from further away. We can take a similar approach for the 5 second limit for the signal to arrive from controller to gate. We can have one event `lower1` which has an activation condition of \perp , and a second `lower2` which has an activation condition of $T_L = 5$, which allows a `lower` event to happen at any point in time up to 5 seconds. Finally, we consider the speed range/non-determinism. There are a few ways to allow a choice from a discrete number of speeds from the range but none of them are ideal. This type of range can be also approached through experimentation on the model by varying parameters. Where non-determinism is used to reason about worst-case scenarios in the context of insufficient knowledge, the goal is to show that the system does not exhibit an unwanted behaviour, no matter what the parameters are (within their prescribed boundaries). In many cases, one can then identify extremal values for quantities of interest to use in the modelling process. For the train gate controller, we can do this by taking the fastest speed as the worst-case analysis, without loss of information and leading to a simpler model.

Thus this section has demonstrated how we can model a standard hybrid systems example in HYPE, as well as illustrating its compositionality both for adding additional components to the model, and adding failure events that depend on variables affected by different subcomponents.

11. Related Work

As mentioned in the introduction, HYPE takes a finer grained, less monolithic approach than the other process algebras for hybrid systems [BM05, vBMR⁺06, RS03, CR05] because it enables the modelling of individual flows. In [Kha06] the comparison of these other process algebras is based on the train gate controller example and in each case, the train, gate and controller components have to be fully, sequentially described and then composed in parallel. This would also be necessary for the modelling of our orbiter example. For each of these process algebras, somewhere in the syntactic description of the system, a term such as $\dot{K} = r_s - r_d - K$, as well as terms for each of the other seven ODEs for the variable K , would need to appear to describe the continuous behaviour that can occur. By comparison, a modeller using HYPE would only need to model the individual flows, and not construct the ODEs explicitly. This could allow non-experts to model hybrid systems more easily.

A classical formalism for expressing hybrid systems is hybrid automata [Hen96]. As presented in Section 4 and Section 7, they are specified by defining explicitly both the control graph and the dynamical conditions within each mode, in terms of differential equations or, more generally, differential inclusions. Two hybrid automata are composed in parallel by synchronising transitions on shared events in the control graph [Hen96]. Flow conditions are combined by taking the logical conjunction of the predicates defining them. Where flows are defined by differential equations, the equation for each variable X must be defined only in one component, otherwise a logical inconsistency may arise. In Corollary 7.1, we showed the difference between the product of two hybrid automata and synchronisation of two HYPE models.

The modelling style of HYPE is quite different. Activities are identified with atomic flows acting on system variables. ODEs are then derived for an individual state by adding the different atomic flows acting on each variable. Activities can change in response to the happening of discrete events, which are controlled not by components but by an external controller and triggered by event conditions. This results in a separation of the description of the response of the system to events from the discrete control structure imposing causality on the happening of events. In contrast to hybrid automata, HYPE allows the separate description of flow conditions, event conditions, and the control graph, making easier the task of modifying the controller or the interactions with the environment. Furthermore, the fact that influence types are defined separately from the structure of the model also separates modelling concerns. Together, these features give a modularity of definition. In addition, compositionality of HYPE manifests on the set of activities (the state of the system) rather than on ODEs, hence we can allow different components of the system to influence the same continuous variable: the combined effect is obtained by superimposing flows, namely by addition on the right hand side of ODEs. Since a HYPE model can be expressed as a hybrid automaton, when using HYPE to model one gets the advantages of HYPE together with the formalism of hybrid automata. This is a distinct advantage over languages such as CHARON [AGLS06], SHIFT [DGV96], and HyCharts [GS02]. These are all compositional formalisms describing hybrid systems which do not map so readily to hybrid automata. They differ from HYPE in that they do not have the simple syntax and structured operational semantics of a process algebra.

Another formalism which has a mechanism to combine flows is hybrid action systems [RRS03]. This language is based on Dijkstra's guarded command language and has predicate transformer semantics. Differential actions consist of guards before ODEs, and parallel composition of these actions is defined as a linear combination of functions with the addition of functions over any shared domains. This differs from our approach where we associate influence names with a specific variable and then sum over all influences for a given variable to obtain the ODE.

Physical systems can also be modelled by constitutive equations and bond graphs [Pay61]. This is a modelling technique in which a system's components are described by means of equations relating main physical quantities of interest. Components are then glued together by imposing suitable conservation laws. This results in a set of differential equations with algebraic constraints, which after an algebraic manipulation, can be simplified by removing variables and constraints. There are also hybrid extensions of the bond graph method which have been represented in a hybrid process algebra [CBM08], to deal with discontinuities in physical systems (such as a bouncing ball).

In HYPE, instead, the modelling activity concentrates around the notion of flow or influence, which is not explicitly connected with physical quantities or with conservation laws, and components are described by specifying the way they react to external events through flows modifications. This may be less natural for certain physical systems, as one has to identify the different influences acting on each variable of interest (without relying on the implicit derivation mechanism provided by conservation laws). However, HYPE's modelling style is straightforwardly applicable to a wider class of systems, at different levels of abstraction.

12. Conclusions and Further Work

In this article, we have presented HYPE, a modelling formalism for hybrid systems, developed its theory and illustrated its applicability. HYPE has novel features that include a fine-grained approach to modelling flows and an explicit controller.

HYPE, in the usual process algebra style, has a syntax consisting of a small number of operators and an operational semantics that describe the behaviour of a HYPE controlled system. This results in a labelled transition system over configurations which are pairs of controlled system derivatives and states. States contain information about the continuous flows that are in operation at a configuration in the labelled transition system, and a transition represents a change to a configuration where different continuous flows apply. This provides a straightforward mapping to a hybrid automaton, where information about the flows in operation are used to construct the ODEs at a mode in the hybrid automaton.

We considered the general behaviour of HYPE models and identified syntactic restrictions called well-definedness and operational restrictions called well-behavedness, to ensure HYPE models have sensible behaviour. We showed for both of these restrictions that they are semantically meaningful in the behaviour they ensure. Moreover, particularly in the case of well-behavedness, the I-graph of the controller with event conditions allows us to reason about the behaviour of a HYPE model without needing to explore it in full detail.

Since hybrid automata are a classical formalism for hybrid systems and also because they provide the hybrid semantics for HYPE models, we considered the expressive differences between products in the two formalisms, showing that HYPE models, due to the use of influences, allow for more expressive and flexible products since the same variable can appear in both models that form the product.

We investigated equivalence semantics and provided a definition of bisimulation over HYPE models together with results about congruence and the ODEs obtained from bisimilar HYPE models. Finally, we presented a classical case study and demonstrated how HYPE can successfully model the problem, as well as an extension of the problem showing the value of compositionality.

In terms of further work, additional modelling using HYPE is important. To date, we have modelled a dual-tank system (as described by [TCT01]), a bottling line (as described by [BM05]), the abstract view of the repressilator [GHB08] (as described by [BP08]) and the circadian clock of a green alga [Gal10] (as described by [AGLT10]).

Since the process algebraic approach uses a language to describe systems, we wish to further explore how reasoning at the language level can provide results about the underlying model. A specific area of focus is that of equivalences that allow us to capture notions of similar behaviour in different systems. Here we could consider bisimulations for models with very different tuples with appropriate maps between variables, events, activities, functions and sets. In the first case these maps would be bijections, but it may be possible

to explore surjective maps. We also wish to investigate links between bisimulation and I-graphs and to obtain results about the well-behavedness of a synchronisation of two HYPE models using the theory developed in Section 6 through constructing an I-graph for the controller $Con_1 \bowtie Con_2$ and the new event conditions ec defined from ec_1 and ec_2 .

We have also extended HYPE models with stochastic durations for events [BGH10a] where the underlying semantic model is Transition-Driven Hybrid Stochastic Automata [BP09], a subset of Piecewise Deterministic Markov Processes [Dav93], and used this extension to model delay-tolerant networks [GHB10, BGH10b]. This is also a direction for future research.

Acknowledgements This research is supported by Royal Society International Joint Project JP090562. Vashti Galpin has been supported by the EPSRC SIGNAL Project, Grant EP/E031439/1. Jane Hillston has been supported by the EPSRC under ARF EP/c543696/01. We thank the reviewers for their comments.

References

- [AGLS06] R. Alur, R. Grosu, I. Lee, and O. Sokolsky. Compositional modeling and refinement for hierarchical hybrid systems. *Journal of Logic and Algebraic Programming*, 68:105–128, 2006.
- [AGLT10] O.E. Akman, M.L. Guerriero, L. Loewe, and C. Troein. Complementary approaches to understanding the plant circadian clock. In *Proceedings of FBTC'10, EPTCS*, 19:1-19, 2010.
- [AHH96] R. Alur, T.A. Henzinger, and P.-H. Ho. Automatic symbolic verification of embedded systems. *IEEE Transactions on Software Engineering*, 22:181–201, 1996.
- [AMP+03] M. Antoniotti, B. Mishra, C. Piazza, A. Policriti, and M. Simeoni. Modeling cellular behavior with hybrid automata: Bisimulation and collapsing. In C. Priami, editor, *Proceedings of CMSB 2003*, LNCS 2602, pages 57–74, 2003.
- [BGH10a] L. Bortolussi, V. Galpin, and J. Hillston. HYPE with stochastic events, 2010. Submitted for publication.
- [BGH10b] L. Bortolussi, V. Galpin, and J. Hillston. Modeling hybrid systems with stochastic events in HYPE. In *Proceedings of the 9th Workshop on Process Algebra and Stochastically Timed Activities (PASTA)*, pages 24–28, 2010.
- [BM05] J.A. Bergstra and C.A. Middelburg. Process algebra for hybrid systems. *Theoretical Computer Science*, 335:215–280, 2005.
- [BP08] L. Bortolussi and A. Policriti. Hybrid approximation of stochastic process algebras for systems biology. In *IFAC World Congress*, Seoul, South Korea, July 2008.
- [BP09] L. Bortolussi and A. Policriti. Hybrid semantics of stochastic programs with dynamic reconfiguration. In *Proceedings of COMPMOD 2009, EPTCS*, 6:63-76, 2009.
- [CBM08] P.J.L. Cuijpers, J.F. Broenink, and P.J. Mosterman. Constitutive hybrid processes: a process-algebraic semantics for hybrid bond graphs. *SIMULATION*, 8:339–358, 2008.
- [CR03] P.J.L. Cuijpers and M.A. Reniers. Hybrid process algebra. Computer Science Reports CSR 03-07, Department of Computer Science, Eindhoven Technical University, 2003.
- [CR05] P.J.L. Cuijpers and M.A. Reniers. Hybrid process algebra. *Journal of Logic and Algebraic Programming*, 62:191–245, 2005.
- [Dav93] M.H.A. Davis. *Markov Models and Optimization*. Chapman & Hall, 1993.
- [DGV96] A. Deshpande, A. Göllü, and P. Varaiya. SHIFT: A formalism and a programming language for dynamic networks of hybrid automata. In P.J. Antsaklis, W. Kohn, A. Nerode, and S. Sastry, editors, *Proceedings of Hybrid Systems IV*, LNCS 1273, pages 113–133, 1996.
- [DPP04] A. Dovier, C. Piazza, and A. Policriti. An efficient algorithm for computing bisimulation. *Theoretical Computer Science*, pages 221–256, 2004.
- [DT07] J.M. Davoren and P. Tabuada. On simulations and bisimulations of general flow systems. In A. Bemporad, A. Bicchi, and G.C. Buttazzo, editors, *Proceedings of HSCC 2007*, LNCS 4416, pages 145–158, 2007.
- [EL00] M.B. Elowitz and S. Leibler. A synthetic oscillatory network of transcriptional regulators. *Nature*, 403:335–338, 2000.
- [Gal10] V. Galpin. Modelling a circadian clock with HYPE. In *Proceedings of the 9th Workshop on Process Algebra and Stochastically Timed Activities (PASTA)*, pages 92–98, 2010.
- [GBH09] V. Galpin, L. Bortolussi, and J. Hillston. Hype: a process algebra for compositional flows and emergent behaviour. In M. Bravetti and G. Zavattaro, editors, *Proceedings of CONCUR 2009*, LNCS 5710, pages 305–320. Springer, 2009.
- [GHB08] V. Galpin, J. Hillston, and L. Bortolussi. HYPE applied to the modelling of hybrid biological systems. *Electronic Notes in Theoretical Computer Science*, 218:33–51, 2008.
- [GHB10] V. Galpin, J. Hillston, and L. Bortolussi. A stochastic hybrid process algebra (poster). In *Models and Logics for Quantitative Analysis (MLQA 2010)*, Edinburgh, July, 2010.
- [GS02] R. Grosu and T. Stauner. Modular and visual specification of hybrid systems: An introduction to HyCharts. *Formal Methods in System Design*, 21:5–38, 2002.
- [Hen96] T.A. Henzinger. The theory of hybrid automata. In *LICS*, pages 278–292, 1996.
- [HH94] T.A. Henzinger and P.-H. Ho. HYTECH: The Cornell HYbrid TECHnology Tool. In P.J. Antsaklis, W. Kohn, A. Nerode, and S. Sastry, editors, *Hybrid Systems*, LNCS 999, pages 265–293. Springer, 1994.

- [Hil05] J. Hillston. Fluid flow approximation of PEPA models. In *Second International Conference on the Quantitative Evaluation of Systems (QEST 2005)*, pages 33–43. IEEE Computer Society, 2005.
- [HKPV95] Thomas A. Henzinger, Peter W. Kopke, Anuj Puri, and Pravin Varaiya. What’s decidable about hybrid automata? In *Proceedings of the twenty-seventh annual ACM Symposium on Theory of Computing, STOC ’95*, pages 373–382, 1995.
- [HTP05] E. Haghverdi, P. Tabuada, and G.J. Pappas. Bisimulation relations for dynamical, control, and hybrid systems. *Theoretical Computer Science*, 342:229–261, 2005.
- [Kha06] U. Khadim. A comparative study of process algebras for hybrid systems. Computer Science Report CSR 06-23, Technische Universiteit Eindhoven, 2006. <http://alexandria.tue.nl/extra1/wskrap/publichtml/200623.pdf>.
- [Kha08] U. Khadim. *Process algebras for hybrid systems: Comparison and development*. PhD thesis, IPA, Technische Universiteit Eindhoven, 2008.
- [LPS00] G. Lafferriere, G.J. Pappas, and S. Sastry. O-minimal hybrid systems. *Mathematics of Control, Signals, and Systems*, 13:1–21, 2000.
- [Mil89] R. Milner. *Communication and concurrency*. Prentice Hall, 1989.
- [MRG05] M.R. Mousavi, M.A. Reniers, and J.F. Groote. Notions of bisimulation and congruence formats for sos with data. *Information and Computation*, 200:107–147, 2005.
- [Pay61] H.M. Paynter. *Analysis and Design of Engineering Systems*. MIT Press, 1961.
- [RRS03] M. Rönkkö, A.P. Ravn, and K. Sere. Hybrid action systems. *Theoretical Computer Science*, 290:937–973, 2003.
- [RS03] W.C. Rounds and H. Song. The ϕ -calculus: A language for distributed control of reconfigurable embedded systems. In O. Maler and A. Pnueli, editors, *Proceedings of HSCC 2003*, LNCS 2623, pages 435–449, 2003.
- [TCT01] B. Tuffin, D.S. Chen, and K.S. Trivedi. Comparison of hybrid systems and fluid stochastic Petri nets. *Discrete Event Dynamic Systems: Theory and Applications*, 11:77–95, 2001.
- [TGH10] M. Tribastone, S. Gilmore, and J. Hillston. Scalable differential analysis of process algebra models. *IEEE Transactions on Software Engineering*, 2010. to appear.
- [vBMR⁺06] D.A. van Beek, K.L. Man, M.A. Reniers, J.E. Rooda, and R.R.H. Schiffelers. Syntax and consistent equation semantics of hybrid χ . *Journal of Logic and Algebraic Programming*, 68:129–210, 2006.
- [vG90] R.J. van Glabbeek. The linear time-branching time spectrum (extended abstract). In *Proceedings of CONCUR 90*, LNCS 458, pages 278–297. Springer, 1990.

A. Proof of Theorem 7.1

Theorem. Let P_1 and P_2 be well-defined HYPE systems $(P_i, \mathcal{V}_i, IN_i, IT_i, \mathcal{E}_i, \mathcal{A}_i, ec_i, iv_i, EC_i, ID_i)$ for $i = 1, 2$. If $\mathcal{V}_1 \cap \mathcal{V}_2 = \emptyset$ then $\mathcal{H}(P_1 \otimes P_2) = \mathcal{H}(P_1) \times \mathcal{H}(P_2)$ when vertices unreachable from the initial vertex are excluded.

Proof. To show that $\mathcal{H}(P_1 \otimes P_2) = \mathcal{H}(P_1) \times \mathcal{H}(P_2)$ we need to show that each component in the two hybrid automata are the same. For the functions that map from vertices and those that map from edges, it is sufficient to show that they map to the same values.

First, consider $P_1 \otimes P_2$ which is the tuple consisting of $(P \stackrel{def}{=} (\Sigma_1 \bowtie \Sigma_2) \bowtie \underline{\text{init.}}(\text{Con}_1 \bowtie \text{Con}_2), \mathcal{V}_1 \cup \mathcal{V}_2, IN_1 \cup IN_2, IT_1 \cup IT_2, \mathcal{E}_1 \cup \mathcal{E}_2, \mathcal{A}_1 \cup \mathcal{A}_2, ec, iv_1 \cup iv_2, EC_1 \cup EC_2, ID_1 \cup ID_2)$ with the conditions and ec given by the definition. We can then construct the labelled transition for this product. Then the hybrid automaton we obtain from the product of the two HYPE models is defined as $H = (V, E, \mathcal{V}_1 \cup \mathcal{V}_2, \mathcal{E}_1 \cup \mathcal{E}_2, flow, init, inv, event, jump, reset, urgent)$ where

- $V = ds((\Sigma_1 \bowtie \Sigma_2) \bowtie \underline{\text{init.}}(\text{Con}_1 \bowtie \text{Con}_2))$
- E contains (v, v') if $v = \langle P, \sigma \rangle$ and $v' = \langle P', \sigma' \rangle$ for some derivation $\langle P, \sigma \rangle \xrightarrow{\underline{a}} \langle P', \sigma' \rangle$
- for $v_j = \langle P_j, \sigma_j \rangle$ then $flow(v_j)[X] = \sum \{r[I(\mathcal{W})] \mid iv(\iota) = X \text{ and } \sigma_j(\iota) = (r, I(\mathcal{W}))\}$
- $init(v) = \begin{cases} res(\underline{\text{init}}), & \text{if } v = \langle P, \sigma \rangle \\ false, & \text{otherwise} \end{cases}$ with primes removed from variables
- $inv(v) = true$.
- Let $e = (\langle P, \sigma \rangle, \langle P', \sigma' \rangle)$ with $\langle P, \sigma \rangle \xrightarrow{\underline{a}} \langle P', \sigma' \rangle$. Then $event(e) = \underline{a}$ and $reset(e) = res(\underline{a})$. Moreover, if $act(\underline{a}) \neq \perp$, then $jump(e) = act(\underline{a})$ and $urgent(e) = true$, otherwise $jump(e) = true$ and $urgent(e) = false$.

Next consider the hybrid automata of P_1 and P_2 . For $i = 1, 2$ we have $H_i = (V_i, E_i, \mathcal{V}_i, \mathcal{E}_i, flow_i, init_i, event_i, jump_i, reset_i, urgent_i)$ where

- $V_i = ds(\Sigma_i \bowtie \underline{\text{init.}} \text{Con}_i)$
- E_i contains (v, v') if $v_1 = \langle P, \sigma \rangle$ and $v' = \langle P', \sigma' \rangle$ for some derivation $\langle P, \sigma \rangle \xrightarrow{\underline{a}} \langle P', \sigma' \rangle$

$flow_i, init_i, inv_i, event_i, reset_i, jump_i$ and $urgent_i$ are defined in the expected manner. The synchronised

product $\mathcal{H}(P_1) \times \mathcal{H}(P_2)$ is $H' = (V_1 \times V_2, E', \mathcal{V}_1 \cup \mathcal{V}_2, \mathcal{E}_1 \cup \mathcal{E}_2, flow', init', event', jump', reset', urgent')$ following the synchronised product definition.

Immediately, we can see that the events and the variables are the same in H and H' . Next we consider the vertices and edges of the graphs that make up the hybrid automata and define a function between these, and show that the function is a bijection, hence we have a graph isomorphism which is sufficient to conclude that the graphs are the same.

Let $W \subseteq V_1 \times V_2$ be the subset of vertices we consider. These are the vertices which are reachable from the initial configuration $(\langle \Sigma_1 \bowtie_* Con_1, \tau_1 \rangle, \langle \Sigma_2 \bowtie_* Con_2, \tau_2 \rangle)$ (τ_1 and τ_2 are the states immediately after init).

We define $f : W \rightarrow V$ by

$$f(\langle \langle \Sigma_1 \bowtie_* D_1, \sigma_1 \rangle, \langle \Sigma_2 \bowtie_* D_2, \sigma_2 \rangle \rangle) = \langle \langle \Sigma_1 \bowtie_* \Sigma_2 \rangle \bowtie_* \langle D_1 \bowtie_* D_2 \rangle, \sigma_1 \cup \sigma_2 \rangle$$

where $D_i \in ds(\text{init}, Con_i)$. Furthermore, for an edge (w, w') , $f(w, w') = (f(w), f(w'))$.

First, we show that if (w, w') is an edge in E' then $f(w, w')$ is an edge in E .

$$\begin{aligned} (w, w') &= ((v_1, v'_1), (v_2, v'_2)) \\ &= (\langle \langle \Sigma_1 \bowtie_* D_1, \sigma_1 \rangle, \langle \Sigma_2 \bowtie_* D_2, \sigma_2 \rangle \rangle, \langle \langle \Sigma_1 \bowtie_* D'_1, \sigma'_1 \rangle, \langle \Sigma_2 \bowtie_* D'_2, \sigma'_2 \rangle \rangle) \end{aligned}$$

$$\begin{aligned} f(w, w') &= (f(w), f(w')) \\ &= (\langle \langle \Sigma_1 \bowtie_* \Sigma_2 \rangle \bowtie_* \langle D_1 \bowtie_* D_2 \rangle, \sigma_1 \cup \sigma_2 \rangle, \langle \langle \Sigma_1 \bowtie_* \Sigma_2 \rangle \bowtie_* \langle D'_1 \bowtie_* D'_2 \rangle, \sigma'_1 \cup \sigma'_2 \rangle) \end{aligned}$$

There are three ways in which an edge can be obtained in a hybrid automata product, and we consider the most complex case only. The other two cases are similar. We consider the case where $event((w, w')) \in \mathcal{E}_1 \cap \mathcal{E}_2$. We can infer that we have two edges $(v_i, v'_i) \in E_1$, $i = 1, 2$ with the same event label as (w, w') . Hence, we have transitions of the form $\langle \Sigma_i \bowtie_* D_i, \sigma_i \rangle \xrightarrow{\underline{a}} \langle \Sigma_i \bowtie_* D'_i, \sigma'_i \rangle$. By manipulation of derivation trees and well-definedness, $\langle \langle \Sigma_1 \bowtie_* \Sigma_2 \rangle \bowtie_* \langle D_1 \bowtie_* D_2 \rangle, \sigma_1 \cup \sigma_2 \rangle \xrightarrow{\underline{a}} \langle \langle \Sigma_1 \bowtie_* \Sigma_2 \rangle \bowtie_* \langle D'_1 \bowtie_* D'_2 \rangle, \sigma'_1 \cup \sigma'_2 \rangle$. Since IN_1 and IN_2 are disjoint, we can start with $\sigma_1 \cup \sigma_2$ and as a result of the derivation, we will obtain $\sigma'_1 \cup \sigma'_2$. We can reverse this argument to show that if $(f(w), f(w'))$ is an edge in E then (w, w') is an edge in E' .

The preceding paragraph implicitly assumes that f is well-defined but it must be proved. For $w \in W$, we show that $f(w) \in V$. Since w is reachable from the initial vertex w_0 , there is a sequence of vertices w_0, w_1, \dots, w_n, w that describe the edges from w_0 to w . Hence, there is a sequence of vertices $f(w_0), f(w_1), \dots, f(w_n), f(w)$. We need to ensure that the initial vertices have the same states and this is guaranteed for well-defined systems by Proposition 5.5.

Next, we need to show that f is a bijection. For $v \in V$, we can use Proposition 5.2 to show that there exists $w \in W$ such that $f(w) = v$. We have indicated above that given an edge $(f(w), f(w')) \in E$ there exist $(w, w') \in E'$. Hence f is surjective. If we have two vertices in W that only differ in one of their states, then the resulting vertices in V will differ, and if they differ in terms of their controller components, then the vertices in V will also differ. If we have two edges in E' , then their images will also differ. Hence f is injective and we have a graph isomorphism.

We need to check the functions on vertices. Clearly, $inv'(w) = inv(f(w))$ since both inv' and inv map to *true*. $init'(w) = false$ for all $w \in W$ except $w' = (\langle \Sigma_1 \bowtie_* Con_1, \tau_1 \rangle, \langle \Sigma_2 \bowtie_* Con_2, \tau_2 \rangle)$ and $f(w') = \langle \langle \Sigma_1 \bowtie_* \Sigma_2 \rangle \bowtie_* \langle Con_1 \bowtie_* Con_2 \rangle, \tau \rangle$. $init((f(w))) = true$ and $init(v) = false$ for all other vertices $v \in V$.

We also need to consider $flow$. For $(v_1, v_2) \in W$, $flow'((v, v)) = flow_1(v_1) \wedge flow_2(v_2)$. Since $V_1 \cap V_2 = \emptyset$, each $flow_i(v_i)$ contributes distinct ODEs on distinct variables to $flow'$ and does not evaluate to *false* for any variable. Consider $flow'(w)[X] = \sum \{r[I(\mathcal{W})] \mid iv_1(\iota) = X \text{ and } \sigma_1(\iota) = (r, I(\mathcal{W}))\}$ where $X \in \mathcal{V}_1$ and $w = (\langle \Sigma_1 \bowtie_* D_1, \sigma_1 \rangle, \langle \Sigma_2 \bowtie_* D_2, \sigma_2 \rangle)$. We know that $f(w) = \langle \langle \Sigma_1 \bowtie_* \Sigma_2 \rangle \bowtie_* \langle D_1 \bowtie_* D_2 \rangle, \sigma_1 \cup \sigma_2 \rangle$ hence $flow(f(w))[X] = \sum \{r[I(\mathcal{W})] \mid (iv_1 \cup iv_2)(\iota) = X \text{ and } (\sigma_1 \cup \sigma_2)(\iota) = (r, I(\mathcal{W}))\}$. Any ι that satisfies $(iv_1 \cup iv_2)(\iota) = X$ must be in IN_1 since $X \in \mathcal{V}_1$, therefore the ODE obtained is the same in both cases because we only need consider σ_1 . An similar argument can be made for $X \in \mathcal{V}_2$.

Next, we need to check the functions on edges. Let $e = (e_1, e_2)$ be an edge in E' . From above, $event'(e) = event(f(e)) = \underline{a}$. For $jump$, we need to consider if $event'(e) = \underline{a}$ is a shared event or not.

- If $\underline{a} \in \mathcal{E}_1 \setminus \mathcal{E}_2$, then there are two cases.
 - If $act_1(\underline{a}) \neq \perp$ then $jump'((e_1, e_2)) = jump_1(e_1) = act_1(\underline{a})$ and $urgent'(e) = urgent_1(e_1) = true$. For H , $jump(f(e)) = act(\underline{a}) = act_1(\underline{a})$ and $urgent(f(e)) = true$ since $act(\underline{a}) \neq \perp$.

- If $act_1(\underline{a}) = \perp$ then $jump'(e) = jump_1(e_1) = true$ and $urgent'(e) = urgent_1(e_1) = false$. For H , $jump(f(e)) = true$ and $urgent(f(e)) = false$.

If $\underline{a} \in \mathcal{E}_2 \setminus \mathcal{E}_1$, the argument is similar.

- If $\underline{a} \in \mathcal{E}_1 \cap \mathcal{E}_2$ then there are again two cases.

- If $act_1(\underline{a}) \wedge act_2(\underline{a}) \neq \perp$, we have $jump'((e_1, e_2)) = jump_1(e_1) \wedge jump_2(e_2) = act_1(\underline{a}) \wedge act_2(\underline{a})$ and $urgent'((e_1, e_2)) = true$. Likewise $jump(f(e)) = act(\underline{a}) = act_1(\underline{a}) \wedge act_2(\underline{a})$ and $urgent(f(e)) = true$ since $act_1(\underline{a}) \wedge act_2(\underline{a}) \neq \perp$.
- Otherwise, $jump'((e_1, e_2)) = jump_1(e_1) \wedge jump_2(e_2) = true$ and $urgent'((e_1, e_2)) = urgent_1(e_1) \vee urgent_2(e_2) = false$, since $act_1(\underline{a}) \wedge act_2(\underline{a}) = \perp$ implies $urgent_1(e_1) = false$ and $urgent_2(e_2) = false$. Likewise, $jump(f(e)) = true$ and $urgent'(f(e)) = false$.

In the case of $reset$ and $reset'$, we need to again consider if we have a shared event or not.

- If $event(e) = \underline{a} \in \mathcal{E}_1 \setminus \mathcal{E}_2$ then $reset'(e) = reset_1(e_1) = res_1(\underline{a})$ and $reset(f(e)) = res(\underline{a}) = res_1(\underline{a})$.
- If $event \underline{a} \in \mathcal{E}_1 \cap \mathcal{E}_2$ then $reset'(e) = reset_1(e_1) \wedge reset_2(e_2) = res_1(\underline{a}) \wedge res_2(\underline{a})$. On the other hand, $reset(f(e)) = res(\underline{a}) = res_1(\underline{a}) \wedge res_2(\underline{a})$.

□