

A Compositional Semantic Theory for Synchronous Component-based Design*

Barry Norton¹, Gerald Lüttgen², and Michael Mendler³

¹ Department of Computer Science, University of Sheffield, UK.
e-mail: b.norton@dcs.shef.ac.uk

² Department of Computer Science, University of York, UK.
e-mail: gerald.luetgen@cs.york.ac.uk

³ Informatics Theory Group, University of Bamberg, Germany.
e-mail: michael.mendler@wiai.uni-bamberg.de

Abstract. Digital signal processing and control (DSPC) tools, such as LabView and iConnect, allow application developers to assemble systems by connecting predefined components in signal-flow graphs and by hierarchically building new components via encapsulating sub-graphs. Run-time environments then dynamically schedule components for execution on some embedded processor, typically in a synchronous cycle-based fashion, and check whether one component jams another by producing outputs faster than can be consumed. Currently, there do not exist formal models of DSPC schedulers that would enable compositional static verification of real-time constraints, such as jam-freeness.

This paper develops a process-algebraic coordination model for synchronous component-based design, which directly lends itself to compositionally formalising the monolithic semantics of DSPC tools. By uniformly combining the well-known concepts of abstract clocks, maximal progress and clock-hiding, it is shown how the DSPC principles of dynamic synchronous scheduling, isochrony and encapsulation may be captured faithfully and compositionally in process algebra, and how observation equivalence may facilitate compositional jam checks. These results provide a foundation for enhancing existing DSPC tools by allowing behavioural validations to be conducted automatically at compile-time.

Keywords: Coordination models, process algebra, digital signal processing, scheduling, jam analysis.

Correspondence: Gerald Lüttgen

Address: Dept. of Computer Science, The University of York
Heslington, York, YO10 5DD, UK

E-mail: gerald.luetgen@cs.york.ac.uk

Phone: +44 190 443-4774

Fax: +44 190 443-2767

* Research supported by EPSRC grant GR/M99637 and British Council grant PRO1163/CH.

1 Introduction

One important domain for embedded-systems designers are *digital signal processing and control applications* (DSPC). These involve dedicated software for control and monitoring problems in industrial production plants, or software embedded in engineering products. The underlying programming style within this domain relies on component-based design. Over many years, engineers have built rich repositories of pre-compiled and well-tested software components (PID-controllers, FIR-filters, FFT-transforms, etc.) that encapsulate technological know-how and hide design complexity behind clear interfaces. Applications can then be programmed efficiently by simply interconnecting components, which frees the application engineer from most of the error-prone low-level programming tasks. Design efficiency is further aided by the fact that DSPC programming tools, including LabView [15], iConnect [25] and Ptolemy [16], typically provide a graphical user interface that supports hierarchical abstraction. Hierarchical extensions of signal-flow graphs permit the *encapsulation* of sub-systems into single components, thus facilitating reuse of system designs.

While the visual signal-flow formalism facilitates mainly the structural design of DSPC applications, the overall behaviour of a component-based system manifests itself only once its components are scheduled and executed on an embedded processor. This *scheduling* is often handled dynamically by run-time environments, as is the case in LabView and iConnect, in order to achieve more efficient and adaptive real-time behaviour as well as some form of control flow. The scheduling typically follows a natural cycle-based synchronous execution model with the phases *collect input* (I), *compute reaction* (R) and *deliver output* (O). This IRO scheduling model is uniformly applied to composite signal-flow graphs as well as their individual components, which may themselves be built hierarchically from smaller entities (cf. Sec. 2). At the top level, the scheduler continuously iterates between executing the *source components* that produce new inputs, e.g., by reading sensor values, and one executing *computation components* that transform input values into output values, which are then delivered to the system environment, e.g., via actuators. Each phase obeys the *synchrony principle* [11], i.e., in (I) all source components are given a chance to collect input from the environment before any computation component is executed, in (R) every computation component whose inputs are available will be scheduled for execution, and in (O) all generated outputs will be delivered before the current cycle ends. The constraint in phase (O), which is known as *isochrony* [12, 18], implies that each output signal will be ‘instantaneously’ received at each connected input. This synchronous scheme can naturally be applied in a hierarchically nested fashion, abstracting a causal sequence of RO-steps produced by a sub-system into a single RO-step.

Like in synchronous programming, the implicit synchrony hypothesis of IRO scheduling assumes that the reaction of a (sub-)system is always faster than its environment issues execution requests. If a component cannot consume its input signals at the pace at which they arrive, a *jam* occurs [25]. In practice, jams usually indicate serious real-time problems (cf. Sec. 2). Unfortunately, in

existing tools, such as iConnect, there are neither compile-time nor run-time checks for detecting jams, thereby forcing engineers to rely on extensive simulations for validating their applications. Moreover, there is no formal model of IRO scheduling for DSPC programming systems that could be used for the compositional analysis of jams, and the question of how to distribute the monolithic IRO scheduler into a uniform coordination model has not been addressed in the literature. Such a model would be extremely useful given that a good deal of real-time validation of DSPC applications could be reduced to jam analysis, in the form of static verification of jam-freeness.

The objective of this paper is to show that a relatively small number of standard concepts studied in concurrency theory provides the key to naturally and *compositionally* formalising the semantics of component-based DSPC designs, and thus to enable static compositional jam checks. The most important concepts from the process-algebra tool-box are *handshake* synchronisation from CCS [19], and *abstract clocks* in combination with *maximal progress* as investigated in temporal process algebras [2], specifically TPL [13], PMC [1] and CSA [6]. We use handshake synchronisation for achieving serialisation, and atomicity and maximal progress clocks for reflecting synchrony. Finally, given maximal progress, synchronous encapsulation may be naturally captured in terms of *clock-hiding*, similar to hiding in CSP [14]. We will uniformly integrate all three concepts into a single process language (cf. Sec. 3), to which we refer as *Calculus for Synchrony and Encapsulation* (CaSE). This calculus conservatively extends CCS in being equipped with a behavioural theory based on observation equivalence [19].

As main contribution, we will formally establish that CaSE is expressive enough for faithfully modelling the principles of IRO scheduling and for reasoning about jams (cf. Sec. 4). First, using a single clock and maximal progress we will show how one may derive a decentralised description of the synchronous scheduler. Second, we prove that isochrony across connections can be modelled via multiple clocks and maximal progress. Third, the subsystems-as-components principle is captured by the clock-hiding operator. Moreover, we will argue that observation equivalence lends itself for statically detecting jams by reducing jam checking to timelock checking.

In the light of these results, our modelling in CaSE yields a *coordination model* for synchronous component-based design, whose virtue is its compositional style for specifying and reasoning about DSPC systems. In particular, our results disprove the perception of designers of DSPC tools that the presence of a global run-time environment and a centralised scheduler precludes the compositional, static capture of semantic properties of DSPC programs, including jam-freeness. Thus, CaSE provides a foundation for developing future-generation DSPC tools that offer the compositional, static analysis techniques desired by engineers.

2 An Example of DSPC Design

Our motivating example is a *digital spectrum analyser*, which is sketched in the signal-flow graph of Fig. 1. The task is to analyse an audio signal and continually

show an array of bar-graphs representing the intensity of the signal in disjoint sections of the frequency range. Our spectrum analyser is designed with help of components Soundcard, Const, Element and BarGraph. Each instance $c1, c2, \dots$ of Element, written as $ck:Element$ or simply ck for $k = 1, 2, \dots$, is responsible for displaying one bar-graph. $ck:Element$ is connected to the single instance $s0$ of component Soundcard, $s0:Soundcard$, which generates the audio signal and provides exactly one audio value each time it is scheduled. It is also connected to instance $sk:Const$ of component Const, which initialises $ck:Element$ by providing filter parameters when it is first scheduled. In contrast to components Soundcard and Const, Element is not a basic but a hierarchical component. Indeed, every ck encapsulates one instance of Filter, Quantise and BarGraph, respectively, namely $ck1:Filter$, $ck2:Quantise$ and $ck3:BarGraph$ as shown in Fig. 2.

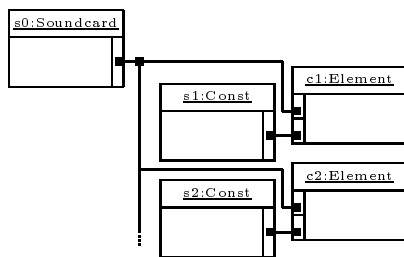


Fig. 1. Example Application




Fig. 2. An instance of Element

Scheduling. According to IRO scheduling, our example application will be serialised as follows within each IRO-cycle. First, each source component instance gets the chance to execute. In the first cycle, this will be $s0:Soundcard$ and all $sk:Const$, which will be interleaved in some arbitrary order. In all subsequent cycles, only $s0:Soundcard$ will request to be scheduled, since $sk:Const$ can only produce a value once. Each produced value will be instantaneously propagated to each $ck:Element$, for all $k \geq 1$, according to the isochronic broadcast. The scheduler then switches to scheduling computation components. Since all necessary inputs of each ck are available in each IRO-cycle, every ck will request to be scheduled. The scheduler will serialise these requests, each ck will execute accordingly, and the current IRO-cycle ends as no outputs generated within ck are to be propagated to its environment. However, since each ck encapsulates further component instances, its execution is non-trivial and involves a sub-scheduler that will schedule $ck1:Filter$, $ck2:Quantise$ and $ck3:BarGraph$ in such a way that an RO-cycle of these instances will appear atomic to ck . This ensures that the scheduling of the inner $ck1$, $ck2$ and $ck3$ will not be interleaved with executing any of the sibling instances cl , for $l \neq k$, of ck .

Isochronic output. Whenever $s0:Soundcard$ is scheduled in our example system, it generates an audio signal whose value is propagated via a wire that forks to port ik of each instance $ck:Element$, for $k \geq 1$. In order for the array of bar-graphs to display a consistent state synchronous with the environment, all ck

must have received the new value from $s0:Soundcard$ before any $cl:Element$ may be scheduled. Thus, $s0:Soundcard$ and all $ck:Element$, for $k \geq 1$, must synchronise to transmit sound values instantaneously. This form of synchronisation is called *isochrony* [12] in hardware, where it is the weakest known synchronisation principle from which non-trivial sequential behaviour can be implemented safely without internal real-time glitches [18].

Jams. Let us now consider what happens if instances $s0:Soundcard$ and $s1:Const$ are accidentally connected the wrong way around, i.e., $s0:Soundcard$ is connected to port $j1$ and $s1:Const$ to port $i1$ of $c1:Element$. Recall that $c11:Filter$ within $c1:Element$ will only read a value, an initialisation value, from port $j1$ in the first IRO-cycle and never again afterwards. Thus, when the value of $s0:Soundcard$ produced in the third cycle is propagated to port $j1$, the system *jams*. This is because the value that has been produced in the second IRO-cycle and stored at this port, has not yet been read by $c11:Filter$. Observe that a jam is different from a deadlock; indeed our example system does not deadlock since all instances of $Element$ other than $c1:Element$ continue to operate properly.

3 CaSE: Calculus for Synchrony and Encapsulation

This section presents our process calculus CaSE, which serves as a framework for deriving our formal coordination model for DSPC design in Sec. 4. CaSE is inspired by Hennessy and Regan's TPL [13], which is an extension of Milner's CCS [19] with regard to syntax and operational semantics. In addition to CCS, TPL includes (i) a *single abstract clock* σ that is interpreted not quantitatively as some number but qualitatively as a recurrent global synchronisation event; (ii) a *timeout operator* $[P]\sigma(Q)$ similar to ATP [20], where the occurrence of σ deactivates process P and activates Q ; (iii) the concept of *maximal progress* [27] that implements the synchrony hypothesis by demanding that a clock can only tick within a process, if the process cannot engage in any internal activity τ .

CaSE further extends TPL by (i) allowing for *multiple clocks* σ, ρ, \dots as in PMC [1] and CSA [6] while, in contrast to PMC and CSA, maintaining the global interpretation of maximal progress; (ii) explicit *timelock* operators Δ and Δ_σ that prohibit the ticking of all clocks and of clock σ , respectively; (iii) *clock-hiding* operators P/σ that internalise all clock ticks of process P . Clock hiding is basically hiding as in CSP [14], i.e., hidden actions are made non-observable. In combination with maximal progress this has the important effect —so far unexplored in the process-algebra community— that all inner clock ticks become included within the synchronous cycle of an outer clock. This is the essence of synchronous encapsulation, as is required by the subsystems-as-components principle in DSPC design. Finally, in contrast to TPL and similar to CCS and CSA, we will equip CaSE with a bisimulation-based semantic theory [19].

Syntax and operational semantics. We let $A = \{a, b, \dots\}$ be a countable set of *input actions* and $\bar{A} = \{\bar{a}, \bar{b}, \dots\}$ be the set of complementing *output actions*. As in CCS [19], an action a communicates with its complement \bar{a} to produce

the *internal action* τ . The symbol \mathcal{A} denotes the set of all actions $\mathcal{A} \cup \overline{\mathcal{A}} \cup \{\tau\}$. Moreover, CaSE is parameterised in a set $\mathcal{T} = \{\sigma, \rho, \dots\}$ of *abstract clocks*, or clocks for brief. The syntax of CaSE is defined by the following BNF:

$$P ::= \mathbf{0} \mid \Delta \mid \Delta_\sigma \mid x \mid \alpha.P \mid P+P \mid P|P \mid P \setminus L \mid P/\sigma \mid [P]\sigma(P) \mid \mu x.P,$$

where x is a *variable* taken from some countably infinite set, and $L \subseteq \mathcal{A} \setminus \{\tau\}$ is a *restriction set*. Further, we use the standard definitions for *static* and *dynamic* operators, *free* and *bound* variables, *open* and *closed* terms, and *guarded* terms. We refer to closed and guarded terms as *processes*, collected in the set \mathcal{P} . For convenience, we write \overline{L} for the set $\{\overline{a} \mid a \in L\}$ and extend the timeout operator to sequences of clocks by defining $[P] =_{\text{df}} P$ and $[P]\sigma_1(Q_1) \dots \sigma_n(Q_n) =_{\text{df}} [[P]\sigma_1(Q_1) \dots \sigma_{n-1}(Q_{n-1})]\sigma_n(Q_n)$. Finally, if P contains actions a_1, a_2, \dots, a_n and the free variable x only, we write $x(a_1, a_2, \dots, a_n) \stackrel{\text{def}}{=} P$ for the process $\mu x.P$. Then, $x(b_1, b_2, \dots, b_n)$ denotes the process $\mu x.P'$, where P' results from P by simultaneously substituting actions a_i by b_i , for $1 \leq i \leq n$.

Table 1. Operational semantics for CaSE

Act	$\frac{}{\alpha.P \xrightarrow{\alpha} P}$	tAct	$\frac{}{\alpha.P \xrightarrow{\alpha} \alpha.P} \alpha \neq \tau$		
Sum1	$\frac{P \xrightarrow{\alpha} P'}{P+Q \xrightarrow{\alpha} P'}$	tNil	$\frac{}{\mathbf{0} \xrightarrow{\alpha} \mathbf{0}}$	tStall	$\frac{}{\Delta_\sigma \xrightarrow{\rho} \Delta_\sigma} \sigma \neq \rho$
Sum2	$\frac{Q \xrightarrow{\alpha} Q'}{P+Q \xrightarrow{\alpha} Q'}$	tSum	$\frac{P \xrightarrow{\alpha} P' \quad Q \xrightarrow{\alpha} Q'}{P+Q \xrightarrow{\alpha} P'+Q'}$		
Res	$\frac{P \setminus L \xrightarrow{\alpha} P' \setminus L}{P \xrightarrow{\alpha} P'} \alpha \notin L \cup \overline{L}$	tRes	$\frac{P \setminus L \xrightarrow{\alpha} P' \setminus L}{P \xrightarrow{\alpha} P'}$		
Par1	$\frac{P Q \xrightarrow{\alpha} P' Q}{Q \xrightarrow{\alpha} Q'}$	tPar	$\frac{P Q \xrightarrow{\alpha} P' Q'}{P \xrightarrow{\alpha} P'} P Q \not\xrightarrow{\tau}$		
Par2	$\frac{P Q \xrightarrow{\alpha} P Q'}{P \xrightarrow{\alpha} P' \quad Q \xrightarrow{\alpha} Q'}$	tHid1	$\frac{P/\sigma \xrightarrow{\tau} P/\sigma}{P \xrightarrow{\rho} P'}$		
Par3	$\frac{P Q \xrightarrow{\tau} P' Q'}{P \xrightarrow{\alpha} P'}$	tHid2	$\frac{P \xrightarrow{\rho} P'}{P/\sigma \xrightarrow{\rho} P'/\sigma} \sigma \neq \rho, P \not\xrightarrow{\tau}$		
Hid	$\frac{P/\sigma \xrightarrow{\alpha} P'/\sigma}{P \xrightarrow{\alpha} P'}$	tTO1	$\frac{}{[P]\sigma(Q) \xrightarrow{\sigma} Q} P \not\xrightarrow{\tau}$		
TO	$\frac{[P]\sigma(Q) \xrightarrow{\alpha} P'}{P[\mu x.P/x] \xrightarrow{\alpha} P'}$	tTO2	$\frac{P \xrightarrow{\rho} P'}{[P]\sigma(Q) \xrightarrow{\rho} P'} \sigma \neq \rho$		
Rec	$\frac{P[\mu x.P/x] \xrightarrow{\alpha} P'}{\mu x.P \xrightarrow{\alpha} P'}$	tRec	$\frac{P[\mu x.P/x] \xrightarrow{\alpha} P'}{\mu x.P \xrightarrow{\alpha} P'}$		

The *operational semantics* of a CaSE process P is given by a labelled transition system $\langle \mathcal{P}, \mathcal{A} \cup \mathcal{T}, \longrightarrow, P \rangle$, where \mathcal{P} is the set of states, $\mathcal{A} \cup \mathcal{T}$ the alphabet, \longrightarrow the transition relation, and P the start state. We refer to transitions with labels in \mathcal{A} as *action transitions* and to those with labels in \mathcal{T} as *clock transitions*. The transition relation $\longrightarrow \subseteq \mathcal{P} \times (\mathcal{A} \cup \mathcal{T}) \times \mathcal{P}$ is defined in Table 1 using

operational rules. For the sake of simplicity, we also write γ for a representative of $\mathcal{A} \cup \mathcal{T}$, as well as $P \xrightarrow{\gamma} P'$ for $\langle P, \gamma, P' \rangle \in \longrightarrow$ and $P \xrightarrow{\gamma}$ for $\exists P' \in \mathcal{P}. P \xrightarrow{\gamma} P'$. Our semantics is set up such that it enjoys a couple of important properties, for all clocks $\sigma \in \mathcal{T}$: (i) *maximal progress*, i.e., $P \xrightarrow{\sigma}$ implies $P \not\xrightarrow{\sigma}$; (ii) *time determinacy*, i.e., $P \xrightarrow{\sigma} P'$ and $P \xrightarrow{\sigma} P''$ implies $P' = P''$. It is time determinacy that distinguishes clock ticks from CSP broadcast communication.

Intuitively, the *nil* process $\mathbf{0}$ permits all clocks to tick, while the *timelock* operators Δ and Δ_σ prohibit the ticking of any clock and of all clocks except σ , respectively. Process $\alpha.P$ may engage in action α and then behave like P . If $\alpha \neq \tau$, it may also idle for each clock σ ; otherwise, all clocks are stopped, thus respecting maximal progress. The *summation operator* $+$ denotes nondeterministic choice, i.e., process $P+Q$ may behave like P or Q . According to time determinacy, time has to proceed equally on both sides of summation, i.e., $P+Q$ can engage in a clock transition and thus delay the nondeterministic choice if and only if both P and Q can. Process $P|Q$ stands for the *parallel composition* of P and Q according to an interleaving semantics with synchronised communication on complementary actions resulting in the internal action τ . Again, time has to proceed equally on both sides of the operator, and the side condition of Rule (tPar) ensures maximal progress. The *restriction operator* $\backslash L$ prohibits the execution of actions in $L \cup \bar{L}$ and thus permits the scoping of actions. The *clock-hiding operator* $/\sigma$ within a process P/σ turns every tick of clock σ in P into the internal action τ . This not only hides clock σ but also pre-empts all other clocks ticking at the same states as σ , according to Rule (tHid2). Process $[P]\sigma(Q)$ behaves as process P , and it can perform a σ -transition to Q , provided P cannot engage in an internal action as is reflected in the side condition of Rule (tTO1). The timeout operator disappears as soon as P engages in some transition labelled differently from σ . Finally, $\mu x.P$ denotes *recursion*, i.e., $\mu x.P$ behaves as a distinguished solution of the equation $x = P$.

Our interpretation of prefixes $\alpha.P$ adopted above, for $\alpha \neq \tau$, is *relaxed* [13], i.e., we allow the process to idle on clock ticks. In the remainder, *insistent prefixes* $\underline{\alpha}.P$ [1], which do not allow clocks to tick, will prove convenient as well. These can be expressed in CaSE by $\underline{\alpha}.P =_{\text{df}} \alpha.P + \Delta$. Similarly, one may define a prefix that only lets clocks not in T tick, for $T \in \mathcal{T}$, by $\underline{\alpha}_T.P =_{\text{df}} \alpha.P + \Delta_T$, where $\Delta_T =_{\text{df}} \sum_{\sigma \in T} \Delta_\sigma$. Finally, we abbreviate $[\Delta]\sigma(P)$ by $\underline{\alpha}.P$.

Temporal observation equivalence and congruence. This section equips CaSE with a bisimulation-based semantics [19]. For the purposes of this paper we will concentrate on *observation equivalence* and *congruence*. The straightforward adaptation of strong bisimulation to our calculus immediately leads to a behavioural congruence, as can easily be verified by inspecting the format of our operational rules and by applying well-known results for structured operational semantics with negative premises [26]. Observational equivalence is a notion of bisimulation in which any sequence of internal τ 's may be skipped. For $\gamma \in \mathcal{A} \cup \mathcal{T}$ we define $\hat{\gamma} =_{\text{df}} \epsilon$ if $\gamma = \tau$ and $\hat{\gamma} =_{\text{df}} \gamma$, otherwise. Further, let $\stackrel{\epsilon}{\xrightarrow{\gamma}} =_{\text{df}} \xrightarrow{\tau}^*$ and $P \stackrel{\epsilon}{\xrightarrow{\gamma}} P'$ if there exist processes P'' and P''' such that $P \xrightarrow{\epsilon} P'' \xrightarrow{\gamma} P''' \xrightarrow{\epsilon} P'$. Carrying over weak bisimulation [19] to CaSE leads to the following definition.

Definition 1. A symmetric relation $\mathcal{R} \subseteq \mathcal{P} \times \mathcal{P}$ is a temporal weak bisimulation if $P \xrightarrow{\gamma} P'$ implies $\exists Q'. Q \xrightarrow{\hat{\gamma}} Q'$ and $\langle P', Q' \rangle \in \mathcal{R}$, for every $\langle P, Q \rangle \in \mathcal{R}$ and $\gamma \in \mathcal{A} \cup \mathcal{T}$. We write $P \approx Q$ if $\langle P, Q \rangle \in \mathcal{R}$ for some temporal weak bisimulation \mathcal{R} .

Temporal observation equivalence \approx is compositional for all operators except summation and timeout. However, for proving compositionality regarding parallel composition and hiding, both of which are defined by operational rules involving negative side conditions, the following proposition is central.

Proposition 1. If $P \approx Q$ and $P \xrightarrow{\sigma} P'$ then $\exists Q', Q''. Q \xrightarrow{\hat{\sigma}} Q'' \xrightarrow{\sigma} Q', P \approx Q'', P' \approx Q'$ and $\{\gamma \in \mathcal{A} \cup \mathcal{T} \mid P \xrightarrow{\gamma}\} = \{\gamma \in \mathcal{A} \cup \mathcal{T} \mid Q'' \xrightarrow{\gamma}\}$.

The validity of this proposition is due to the maximal-progress property in CaSE. This is also why we do not need to equip temporal observation equivalence with complex conditions on initial action sets, as is necessary in calculi incorporating a weaker notion of maximal progress [6]. As usual, observation equivalence is not compositional for the choice operators summation and timeout. To identify the largest equivalence contained in \approx , the summation fix of CCS is not sufficient. As in other work in temporal process algebras [27], the deterministic nature of clocks implies the following definition of *temporal observation congruence*.

Definition 2. A symmetric relation $\mathcal{R} \subseteq \mathcal{P} \times \mathcal{P}$ is a temporal observation congruence if for every $\langle P, Q \rangle \in \mathcal{R}$, $\alpha \in \mathcal{A}$ and $\sigma \in \mathcal{T}$:

1. $P \xrightarrow{\alpha} P'$ implies $\exists Q'. Q \xrightarrow{\hat{\alpha}} Q'$ and $P' \approx Q'$.
2. $P \xrightarrow{\sigma} P'$ implies $\exists Q'. Q \xrightarrow{\hat{\sigma}} Q'$ and $\langle P', Q' \rangle \in \mathcal{R}$.

We write $P \cong Q$ if $\langle P, Q \rangle \in \mathcal{R}$ for some temporal observational congruence \mathcal{R} .

As desired, we obtain the following result, whose proof is standard [6, 19].

Thm 1. The equivalence \cong is the largest congruence contained in \approx .

It is not difficult to see that CaSE is a conservative extension of CCS [19]. Indeed, CCS can be identified in terms of syntax, operational semantics and bisimulation semantics as the sub-calculus of CaSE which is obtained by defining $\mathcal{T} = \emptyset$. For finite-state systems, temporal observational equivalence can be computed efficiently, using standard partition-refinement algorithms as implemented in existing verification tools, such as the CWB-NC [7].

4 A Synchronous Coordination Model with Encapsulation

This section introduces our coordination model for DSCP applications on the basis of our process calculus CaSE. Particular emphasis is given regarding the issues of component instantiation, synchronous scheduling, isochronic forks, and jam analysis. We illustrate our modelling using the digital-spectrum-analyser example introduced in Sec. 2.

We start off with attaching behavioural descriptions to basic components, which describe their interface behaviour to a scheduler. Informally, we say that the states of component descriptions must each belong to one of the classes ‘input’, ‘ready’, ‘waiting’, ‘internal’, ‘output’ and ‘finished’. Input states may be non-deterministic and are source of transitions labelled from a set $I \subset \Lambda$ to either input states or ready states. Ready states must have one transition labelled \bar{r} , a request to execute, to a waiting state which must have one transition labelled g , which is a signal that the request is granted, leading to an internal state. Internal states again have non-deterministic transitions, but labelled only in the silent action τ , representing the progress and resulting state of the internal computation; that this may vary, being dependent on the values of the input data, is modelled by non-determinism. The destination of internal states must be output states from which a deterministic sequence $\bar{o}_1, \dots, \bar{o}_n$ of outputs, i.e., members of the set $O \subset \Lambda$, is produced by transitions to output states or a finished state, at which point all output is delivered. The sequence $\vec{o} = [\bar{o}_1, \dots, \bar{o}_n]$ is chosen non-deterministically by the internal τ -computation that leads into the output region. Finished states have one transition labelled \bar{g} , signifying that the thread of control is being handed back, to either an input state or $\mathbf{0}$, meaning that the component is finished. Observe the scheduling sequence $r \cdots g \cdots \bar{g}$, in which the component requests execution by r , obtains the execution grant through g and finally signals completion with \bar{g} .




Fig. 3. Fig. 4. Fig. 5. Fig. 6. Fig. 7.

For any model, I, O and $\{r, g\}$ must be disjoint and in the following we use only a, b and decorations thereof to name input channels, and range over these with i , and c, d and decorations thereof to name outputs, ranged over by o . Having defined as usual the ‘syntactic sort’ of a process $\mathcal{S}(P) \subset \Lambda$, we define also the ‘input sort’ and ‘output sort’, $\mathcal{I}(P) \subset I, \mathcal{O}(P) \subset O$, in the natural way. Note that the request and grant ports $\{r, g\}$, which connect a component with its scheduling environment, are also part of the component sort. Moreover, we may use the following parameterised definitions to define the typical component behaviour that is consistent with the above interface description, such as those seen in Figs. 3–7 for our example application.

$$\begin{aligned}
 I(\vec{i}, P) &\stackrel{\text{def}}{=} \begin{cases} \sum_{1 \leq j \leq n = |\vec{i}|} i_j \cdot I([i_1, i_2, \dots, i_{j-1}, i_{j+1}, \dots, a_n], P) & \text{if } |\vec{i}| > 0 \\ \bar{r}.g.P & \text{otherwise} \end{cases} \\
 O(\vec{o}, P) &\stackrel{\text{def}}{=} \begin{cases} \bar{o}_1 \cdot O([o_2, \dots, o_{|\vec{o}|}], P) & \text{if } |\vec{o}| > 0 \\ \bar{g}.P & \text{otherwise} \end{cases} \\
 C(\vec{i}, \vec{o}) &\stackrel{\text{def}}{=} I(\vec{i}, \sum_{1 \leq j \leq |\vec{o}|} \tau \cdot O(\vec{o}_j, C(\vec{i}, \vec{o})))
 \end{aligned}$$

First, $I(\vec{i}, P)$ defines an input phase in the vector of inputs \vec{i} which continues, when scheduled, in the internal state P ; $O(\vec{o}, P)$ defines an output phase with fixed sequence of outputs $\vec{o} = [\vec{o}_1, \dots, \vec{o}_n]$, continuing in input state P having handed back the thread of control; $C(\vec{i}, \vec{o})$ defines a typical cyclic component in its input state with input ports \vec{i} and a choice of possible output sequences represented as a (finite) sequence of sequences \vec{o} of individual output actions. Each component \vec{o}_j , for $1 \leq j \leq |\vec{o}|$, of \vec{o} may be an arbitrary (finite) sequence $\vec{o}_j = [\vec{o}_{j1}, \dots, \vec{o}_{jn_j}]$ of output actions, $o_{ji} \in O$. In this way we can model arbitrary output behaviour. Each execution cycle of $C(\vec{i}, \vec{o})$ involves all inputs from \vec{i} and non-deterministic output behaviour defined by the members of \vec{o} .

Returning to our example, observe that the CaSE process of Filter given in Fig. 5, $I(\{a, b\}, \tau.\bar{c}.\bar{g}.C([a], [[c]]))$, features the desired behaviour as described in Sec. 2. In particular, it reads from *both* inputs a, b only in the first scheduling cycle. In all subsequent ones, i.e., after passing once through $\bar{r}.g.\tau.\bar{c}.\bar{g}$, it behaves as the cyclic process $C([a], [[c]])$ that only ever consumes from input a .

4.1 Component instantiation. A pure component as introduced above uses the $\{r, g\}$ interface to negotiate its execution with its environment. From the point of view of the component, it does not matter whether it communicates with a centralised or a distributed scheduler. In this section, we develop a concept of *wrappers* for harnessing components with enough local control so they participate coherently in a global IRO scheduling scheme, without presence of a global scheduler. Indeed all wrappers added together represent a distributed version of an imagined central IRO scheduler. Wrapping a component is an instantiation into a communication *and* scheduling discipline, as in Ptolemy [16]. Our general model for a ‘wrapper’ is the following:

$$\begin{aligned} \text{InstWrapper}(\text{SchedIdiom}, \text{CommIdiom}, \sigma_s)[-] &\stackrel{\text{def}}{=} \\ &((- | \text{SchedIdiom}(\sigma_s, -)) \setminus \mathcal{O}(-) \cup \{r, g\} | \\ &\quad \Pi_{o \in \mathcal{O}(-)} \text{CommIdiom}(o)) \setminus \{b_o, f_o \mid o \in \mathcal{O}(-)\}, \end{aligned}$$

where SchedIdiom and CommIdiom are scheduling and communication idiom, respectively. To implement IRO scheduling we consider two variants of SchedIdiom , namely CompInst and SourceInst to be discussed below.

Synchronous scheduling. Given a basic component C , we can instantiate it either as a computation component, using idiom $\text{CompInst}(\sigma_s, C)$, or as a source component with $\text{SourceInst}(\sigma_s, C)$, where σ_s represents the phase clock. This clock organises strict alternation between source and computation phases and, by way of maximal progress, implements run-to-completion within each phase. To achieve serialisation on a single thread of control, a token-passing style is used, where a component may only execute if it possesses the execution token and surrenders this when the execution is complete.

The token is passed on label t_c between computation components, and on label t_e between source components. Furthermore, each output event o is split systematically into a request-acknowledge pair $\bar{b}_o.f_o$, to prepare for isochronic

output distribution:

$$\begin{aligned} \text{CompInst}(\sigma_s, -) &\stackrel{\text{def}}{=} r \cdot (\bar{r}_e \cdot t_c \cdot \bar{g} \cdot \text{Inst}(t_c, t_e, \text{CompInst}(\sigma_s, -), \sigma_s, -) + \\ &\quad t_c \cdot \bar{g} \cdot \text{Inst}(t_c, t_e, \text{CompInst}(\sigma_s, -), \sigma_s, -)) \\ \text{SourceInst}(\sigma_s, -) &\stackrel{\text{def}}{=} \sigma_s \cdot r \cdot t_e \cdot \bar{g} \cdot \text{Inst}(t_e, t_c, \text{SourceInst}(\sigma_s, -), \sigma_s, -) \\ \text{Inst}(t_1, t_2, P, \sigma_s, -) &\stackrel{\text{def}}{=} \sum_{o \in \mathcal{O}(-)} o \cdot \bar{b}_o \cdot f_o \cdot \text{Inst}(t_1, t_2, P, \sigma_s, -) + g \cdot [\bar{t}_1 \cdot P] \sigma_s (\bar{t}_2 \cdot P) \end{aligned}$$

The idiom $\text{CompInst}(\sigma_s, C)$ will accept the request signal r from component C and wait for an execution token t_c from some other instantiated peer component. When there is no execution token present at the same hierarchy level, $\text{CompInst}(\sigma_s, C)$ instead will issue a request \bar{r}_e to its environment and then wait for t_c . Once t_c has arrived, $\text{CompInst}(\sigma_s, C)$ will grant C its original request with signal g . The wrapper then behaves as $\text{Inst}(t_c, t_e, \text{CompInst}(\sigma_s, C), \sigma_s, C)$ while component C will execute. In this state, whenever C finishes and has data to distribute, the scheduling idiom passes on each output signal o to the communication idiom as b_o (understood as a ‘broadcast o ’ signal) and waits for a signal f_o (understood as ‘finished broadcasting o ’) in return. Once the component cycle is finished, signalled by a return of grant g , an attempt is made to pass on the token $t_1 = t_c$ to one of the peers. If, however, even in the presence of this offer, clock σ_s ticks thereby signalling the end of the phase, then the other token $t_2 = t_e$ is passed out, which will start the source phase. The source idiom $\text{SourceInst}(\sigma_s, C)$ is analogous to $\text{CompInst}(\sigma_s, C)$, except that t_e and t_c are interchanged. Also, since sources must execute at most once per cycle, their idiom $\text{SourceInst}(\sigma_s, C)$ starts with an initial phase clock. In this way the clock has to tick between any two executions of the same source component.

Consider the Filter component of our example system of Fig. 2, whose behavioural definition in Fig. 5 we also abbreviate as Filter. According to our wrapper definition, the first step in instantiating this computation component is as follows:

$$\text{Filter}' \stackrel{\text{def}}{=} (\text{Filter} \mid \text{SchedIdiom}(\sigma_s, \text{Filter})) \setminus \{c, r, g\},$$

with the new ‘external’ interface sort consisting of $\{t_c, t_e, r_e, \sigma_s\}$ for the scheduling and $\{a, b, b_c, f_c\}$ for data input and output.

Isochronic Broadcast. There are two parts to form compositional isochronic forks in our coordination model. The first is a ‘broadcast agent’ $\text{IsoBroad}(o)$, which has been referred to as $\text{CommIdiom}(o)$ above:

$$\begin{aligned} \text{IsoBroad}(o) &\stackrel{\text{def}}{=} \underline{b}_{o\sigma_o} \cdot \text{IsoBroad}'(o) \\ \text{IsoBroad}'(o) &\stackrel{\text{def}}{=} [\bar{\sigma} \cdot \text{IsoBroad}'(o)] \sigma_o (\bar{f}_o \cdot \text{IsoBroad}(o)) \end{aligned}$$

For each broadcast request b_o of the component wrapped, an arbitrary number of ‘copies’ of each signal will be communicated on $\bar{\sigma}$ until the clock σ_o defining

the isochronous instant in which the communication occurs ticks and ends that instant. Because of maximal progress σ_o can only proceed when there are no further receivers listening on o . In this way the signal o obtains maximal distribution. Only when all receivers are saturated will \bar{f}_o occur, thereby signalling the run-to-completion of the broadcast back to the component. Note that isochrony cannot be modelled faithfully in Hoare's CSP [14] or Prasad's CBS [23]. While the broadcasting primitive in CSP does not distinguish between senders and receivers and thus ignores the direction in which information is propagated, the one in CBS does not force receivers to synchronise with senders.

We can now complete our instantiation of the Filter component, extending Filter' by bolting on one IsoBroad(c) for the output line c . This gives

$$\text{Filter}'' \stackrel{\text{def}}{=} (\text{Filter}' \mid \text{IsoBroad}(c)) \setminus \{b_c, f_c\}$$

which is the same as InstWrapper(CompInst, IsoBroad, σ_s)[Filter]. Note that the scheduling interface of Filter now is $\{t_c, t_e, r_e, \sigma_s, \sigma_c\}$, which is extended by the isochrony clock σ_c and for data signalling $\{a, b, c\}$. The output broadcast which was controlled in Filter' by the pair b_c, f_c is now handled by the single line c together with the clock σ_c . The fact that control is now via a clock is what will make the wiring up of Filter'' with an arbitrary number of broadcast receivers compositional.

For the following, we assume that we have obtained computation component instances Quantise'' and BarGraph'' in a similar way. These can then be used to assemble the sub-systems Element, as well as source components Soundcard'', Const'' required for the example application of Fig.1. All these are InstWrapper instantiations of the corresponding basic components seen in Figs. 3–7.

4.2 Isochronic wiring. Now that we have instantiations of our components we need to wire them up. In our setting wires are specific agents that actively participate in broadcasts along them. Our 'wire agents'

$$\text{IsoWire}(o, \sigma_s, i) \stackrel{\text{def}}{=} o . \bar{i}_{\{\sigma_s, \sigma_o\}} . \sigma_o . \text{IsoWire}(o, \sigma_s, i)$$

connect a producer on output o with a consumer on input i . Whenever this agent IsoWire(o, σ_s, i) picks up a signal o it blocks clock σ_o until it has successfully delivered with \bar{i} . Then, σ_o must tick for the wire to cycle back and be ready again. In this way, a single transmission along the wire is sandwiched between o and σ_o . Since σ_o is a global event, we can compositionally join together and synchronise an arbitrary number of parallel wire transmissions. To formalise this, we introduce *closed forks*, defined in the CCS sub-calculus of CaSE:

$$\begin{aligned} \text{Fork}(o, \vec{i}) &\stackrel{\text{def}}{=} b_o . \text{Fork}'(o, \vec{i}, \vec{i}) \\ \text{Fork}'(o, \vec{i}, \vec{j}) &\stackrel{\text{def}}{=} \begin{cases} \sum_{l \leq n = |\vec{j}|} \bar{j}_l . \text{Fork}'(o, \vec{i}, [j_1, j_2, \dots, j_{l-1}, j_{l+1}, \dots, j_n]) & \text{if } |\vec{j}| > 1 \\ \bar{j}_1 . \bar{f}_o . \text{Fork}(o, \vec{i}) & \text{if } |\vec{j}| = 1 \end{cases} \end{aligned}$$

The following theorem shows that our isochronic forks behave equivalently to closed forks, for any number of processes attached to the isochronic broadcast.

Thm 2. $(IsoBroad(o) | \Pi_{j \leq |\vec{i}|} IsoWire(o, \sigma_s, i_j)) / \sigma_o | \Delta_{\sigma_s} \cong Fork(o, \vec{i}) | \Delta_{\{\sigma_s, \sigma_o\}}$

Note that in our modelling, the synchronous phase clock σ_s is admitted only in the initial state, i.e., as long as the wire is 'empty'. The wire's ability to pass on its previous value and thus to become empty, however, depends on the input-readiness of the component instances to which \vec{i} connects forward. If any such instance is not ready, clock σ_s will never again tick, turning this local 'jam' condition into a timing flaw, which is a global condition.

Returning to our example, we could build an instantiated sub-system for Element, which overall behaves like a computation component, as follows:

$$\begin{aligned} \text{Subsystem}'' \stackrel{\text{def}}{=} & (\text{Filter}''(c, d, o_1) | \text{Isowire}(o_1, \sigma_s, i_1) | \text{Quantise}''(i_1, o_2) | \\ & \text{Isowire}(o_2, \sigma_s, i_2) | \text{BarGraph}''(i_2) \\ &) \setminus \{o_1, i_1, o_2, i_2\} / \{\sigma_{o_1}, \sigma_{o_2}\}, \end{aligned}$$

where c, d, o_1, o_2, i_1, i_2 are some arbitrarily chosen names for the input and output ports of the components. Note that we have restricted away the internal channels $\{o_1, i_1, o_2, i_2\}$ and hidden the internal isochrony clocks $\{\sigma_{o_1}, \sigma_{o_2}\}$. For the top-level system of Fig. 1 a similar composition, say *Application''*, can be formed from one source instance *Soundcard''* and a suitable number of source instances *Const''* and component instances *Element''*. We discuss next how one can obtain *Element''* from a composite sub-system such as *Subsystem''* by encapsulation.

4.3 Encapsulation. To encapsulate a signal-flow graph such as *Subsystem''* as a basic component in itself, we finally define another 'wrapper', which is in some sense the inverse of the instantiation wrapper:

$$\begin{aligned} \text{EncWrapper}(\sigma_s)[_] \stackrel{\text{def}}{=} & (_ | \text{Enc}(\sigma_s)) \setminus \{t_c, t_e, r_e\} / \sigma_s \\ \text{Enc}(\sigma_s) \stackrel{\text{def}}{=} & \underline{r_e} \cdot \bar{r} \cdot \underline{g} \cdot \bar{t}_c_{\sigma_s} \cdot t_e \cdot \bar{g}_{\sigma_s} \cdot \text{Enc} \quad . \end{aligned}$$

The *EncWrapper* translates back the scheduling interface $\{t_c, t_e, r_e, \sigma_s\}$ into $\{r, g\}$, which is the interface of a pure component, while keeping the signal input and output ports intact. In the parallel composition *Subsystem''* | *Enc*(σ_s), the process *Enc*(σ_s) picks up any request for token r_e from the computation components inside *Subsystem''*, passes it out as a request r , then waits for the grant signal g , upon which gives the execution token down into *Subsystem''* via \bar{t}_c . At that point, it waits patiently for signal t_e from *Subsystem''*, which indicates that *Subsystem''* has finished one phase cycle. Then, *Enc*(σ_s) finishes off by handing out the \bar{g} signal to its environment, whence signalling completion of one computation run. Thus, seen from the outside, *Subsystem''* | *Enc*(σ_s) behaves like a basic component, while *Enc*(σ_s) emulates a token-passing environment to the inner *Subsystem''*.

EncWrapper is obtained from *Subsystem''* | *Enc*(σ_s) by restricting away all token passing signals $\{t_c, t_e, r_e\}$ and hiding the internal synchronous phase clock σ_s . All ticks of σ_s are turned into τ 's, which from the point of view of the environment now count as proper internal computation steps. The resulting encap-

ulated system $\text{Element} \stackrel{\text{def}}{=} \text{EncWrapper}(\sigma_s)[\text{Subsystem}'']$ is a fresh basic component. It may be instantiated again as a computation component using InstWrapper to give $\text{Element}'' \stackrel{\text{def}}{=} \text{InstWrapper}(\text{CompInst}, \text{IsoBroad}, \sigma_s)[\text{Element}]$, which may be assembled into the complete system $\text{Application}''$, as suggested above.

Note that a signal-flow graph to be encapsulated may only sensibly be made up of instances of computation components and the new component will be a computation component in any inputs not restricted away. The values communicated by all wires should also first be restricted, as should the values offered by broadcast outputs and the clocks that bound them should be hidden. In order to form an output from the new component, a wire should be connected to supply the value at a desired port whose name is not then restricted. Bearing this in mind, we can now formally state the desired encapsulation property, namely that the encapsulation of a component instance should behave equivalently to the original component.

Thm 3. *Let Comp be an arbitrary component. Then*

$$\begin{aligned} & \text{EncWrapper}(\sigma_s)[(\text{InstWrapper}(\text{CompInst}, \text{IsoBroad}, \sigma_s)[\text{Comp}] \\ & \quad | \prod_{o \in \mathcal{O}(\text{Comp})} \text{IsoWire}(o, \sigma_s, o)) / \vec{\sigma}] \\ & \cong \text{Comp} | \Delta_{\{\sigma_s\} \cup \vec{\sigma}}, \end{aligned}$$

where $\vec{\sigma}$ consists of all isochronic clocks σ_o , for $o \in \mathcal{O}(\text{Comp})$.

4.4 Jams analysis. As suggested above, a jam is said to occur when a ‘wire’ is unable to pass on the value it is holding. In our model this produces a path to a state where $\bar{i}.P + \Delta_{\{\sigma_s, \sigma_{\text{name}}\}}$ is a *timelocked* parallel component in the global state. Since the isochronic broadcast agent cannot complete until its clock ticks, it will not signal completion and so the relevant instantiation wrapper will not release the token and the system is deadlocked. Thus, a jam manifests itself as a timelock within our compositional coordination model. This is made explicit in the following theorem.

Thm 4. *If System possesses only τ - and σ_s -transitions and no infinite τ -computations, then the following holds, where $\text{Check} \stackrel{\text{def}}{=} \mu x. [\Delta] \sigma_s(x)$:*

$$\text{System} \approx \text{Check} \quad \text{iff} \quad \nexists s \in \{\tau, \sigma_s\}^*. \text{System} \xrightarrow{s} P \not\stackrel{q_s}{\rightarrow} .$$

We conclude this section with two observations on how time-lock and thus jams may be escaped. Further parallel composition of a jammed system may lead to the ability to escape the insistent states causing the jams via communication, which is exactly how an engineer is supposed to resolve the problem. Alternatively, the scheduling clock may be hidden, but this may only be done in our coordination model via encapsulation. It is undesirable for jams to be hidden by encapsulation, as they are design faults that should be revealed. Indeed, in our planned integration of our coordination model within a type system for DSPC programming environments, which will be expanded upon in the next section, encapsulating faulty components will be disallowed.

5 Related Work

To the best of our knowledge, our coordination model is the first formal model of the synchronous and hierarchical scheduling discipline behind DSPC programming systems. Our process–algebraic approach complements existing work in *distributed object-oriented systems* [21] and in *architectural description languages* [3]. There, the focus is on distributed software rather than embedded centralised systems, and consequently on asynchronous rather than synchronous component behaviour. However, in both application domains, object-oriented systems and software architectures, formalisms have been investigated for describing and reasoning about the interface behaviour of components, too.

In object-oriented systems, simple automata-based frameworks have been studied, where finite automata model the life-cycle of objects [21]. Within these frameworks, one may then reason at compile-time whether each invocation of an object’s method at run-time is permissible. This semantic analysis is different from jam analysis in DSPC applications, but similar to the *compatibility analysis* of interface automata employed in Ptolemy [9], which we will discuss below. A process–algebraic model of this theory for object-oriented design has been developed as well and is presented in [24]. In architectural description languages, the formalism of process algebra has been studied by Bernardo et al. [3]. Their approach rests on the use of CSP-style broadcast communication together with asynchronous parallel composition. Like in our application domain of DSPC design, the intention is to identify communication problems, but these are diagnosed in terms of deadlock behaviour [4].

As illustrated earlier, deadlock is a more specific property than the jam property investigated by us: a jam in one component jams the whole system, but a deadlock in one component does not necessarily result in a system deadlock. To observe ‘local’ deadlocks in single components, a theory of *location equivalence* [5] has been developed in the literature, which refines Milner’s theory of observation equivalence [19] and observes the location, or system component, from which an action is performed. Observation equivalence is unnecessarily expressive and complicated for the purposes of static jam analysis. Indeed, we have shown that there is no need to refer to locations when analysing jams, for which one may simply check via temporal observation equivalence.

From a practical point of view, we envision our coordination model based on the process calculus CaSE to play the role of a *reactive-types* language [22]. This would enable designers to specify the intended interactions between a given component and its environment as a type, and permit tool implementations to reduce type checking to temporal observation–equivalence checking. This idea is somewhat similar to the one of *behavioural types* in the Ptolemy community [17]. Behavioural types in Ptolemy are based on the formalism of *interface automata* [8] and employed for checking the so-called *compatibility* property between components [9]. However, interface automata are not expressive enough to reason about jams, which Ptolemy, for the restricted class of synchronous data-flow (SDF) models, handles by linear–algebra techniques. In contrast, CaSE’s semantic theory is more general than SDF and lends itself to compositionally

checking jams at compile-time. Note in this context that our proposed notion of reactive type is different from the one studied for the synchronous language Signal [10]. Behavioural types in Signal are defined in terms of Signal's clock calculus, which focuses on data-flow rather than signal-flow.

6 Conclusions and Future Work

This paper presented a novel compositional coordination model for the synchronous component-based design of and reasoning about DSPC applications, thus complementing work in distributed object-oriented systems and architectural description languages. Our coordination model benefited from several ideas that have been investigated in the field of concurrency theory. In particular, we demonstrated that the semantic concepts underlying the IRO principle of DSPC tools, namely dynamic synchronous scheduling, isochrony and encapsulation, can be captured by uniformly combining the process-algebraic concepts of abstract clocks, maximal progress and clock hiding. To do so, we defined the process calculus CaSE and developed its behavioural theory based on temporal observation equivalence. This equivalence was then used to prove that clocks and maximal progress are indeed sufficient for compositionally describing DSP schedulers, including the isochronic propagation of output signals, and that clock hiding reflects the encapsulation process in hierarchical design. In addition, CaSE facilitates the static, compositional reasoning about jams in DSPC applications, via observation-equivalence checks against timelocked processes.

Future work should proceed along two orthogonal directions. First, CaSE should be integrated in DSPC tools in the form of a reactive-types system. Second, the semantic theory of CaSE should be completed by providing an axiomatisation of temporal observation congruence for regular processes.

References

1. H.R. Andersen and M. Mendler. An asynchronous process algebra with multiple clocks. In *ESOP '94*, vol. 788 of *LNCS*, pp. 58–73, 1994.
2. J.C.M. Baeten and C.A. Middelburg. *Process Algebra with Timing*. EATCS Monograph. Springer-Verlag, 2002.
3. M. Bernardo, P. Ciancarini, and L. Donatiello. On the formalisation of architectural types with process algebras. In *MASCOTS 2000*, pp. 140–148. ACM Press, 2000.
4. M. Bernardo, P. Ciancarini, and L. Donatiello. Detecting architectural mismatches in process algebraic descriptions of software systems. In *WICSA 2001*, pp. 77–86. IEEE Comp. Soc. Press, 2001.
5. G. Boudol, I. Castellani, M. Hennessy, and A. Kiehn. Observing localities. *TCS*, 114(1):31–61, 1993.
6. R. Cleaveland, G. Lüttgen, and M. Mendler. An algebraic theory of multiple clocks. In *CONCUR '97*, vol. 1243 of *LNCS*, pp. 166–180, 1997.
7. R. Cleaveland and S. Sims. The NCSU Concurrency Workbench. In *CAV '96*, vol. 1102 of *LNCS*, pp. 394–397, 1996.

8. L. de Alfaro and T.A. Henzinger. Interface automata. In *ESEC/FSE 2001*, vol. 26, 5 of *Softw. Eng. Notes*, pp. 109–120. ACM Press, 2001.
9. A. Chakrabarti et al. Interface compatibility checking for software modules. In *CAV 2002*, vol. 2404 of *LNCS*, pp. 428–441, 2002.
10. P. le Guernic, T. Gautier, M. le Borgne, and C. le Maire. Programming real-time applications with Signal. *Proc. of the IEEE*, 79(9):1305–1320, 1991.
11. N. Halbwachs. *Synchronous Programming of Reactive Systems*. Kluwer, 1992.
12. S. Hauck. Asynchronous design methodologies: An overview. *Proc. of the IEEE*, 83(1):69–93, 1995.
13. M. Hennessy and T. Regan. A process algebra for timed systems. *Inform. and Comp.*, 117:221–239, 1995.
14. C.A.R. Hoare. *Communicating Sequential Processes*. Prentice Hall, 1985.
15. G.W. Johnson and R. Jennings. *LabVIEW Graphical Programming*. McGraw-Hill, 2001.
16. E.A. Lee. Overview of the Ptolemy project. Techn. Rep. UCB/ERL M01/11, Univ. of California at Berkeley, 2001.
17. E.A. Lee and Y. Xiong. Behavioral types for component-based design. Techn. Rep. UCB/ERL M02/29, Univ. of California at Berkeley, 2002.
18. A.J. Martin. The limitations to delay-insensitivity in asynchronous circuits. In *6th MIT Conf. on Advanced Research in VLSI*, pp. 263–287. MIT Press, 1990.
19. R. Milner. *Communication and Concurrency*. Prentice Hall, 1989.
20. X. Nicollin and J. Sifakis. The algebra of timed processes, ATP: Theory and application. *Inform. and Comp.*, 114:131–178, 1994.
21. O. Nierstrasz. Regular types for active objects. In *OOPSLA '93*, ACM SIGPLAN Notices, pp. 1–15, 1993.
22. B. Norton. Reactive types for dataflow-oriented component-based development. In *AVoCS 2002*, vol. CSR-02-6 of *Techn. Report Series*. Univ. of Birmingham, 2002.
23. K.V.S. Prasad. Programming with broadcasts. In *CONCUR 93*, vol. 715 of *LNCS*, pp. 173–187, 1993.
24. F. Puntigam. Type specifications with processes. In *FORTE '95*, vol. 43 of *IFIP Conf. Proc.* Chapman & Hall, 1995.
25. A. Sicheneder et al. Tool-supported software design and program execution for signal processing applications using modular software components. In *STTT '98*, BRICS Notes Series NS-98-4, pp. 61–70, 1998.
26. C. Verhoef. A congruence theorem for structured operational semantics with predicates and negative premises. *Nordic J. of Computing*, 2(2):274–302, 1995.
27. W. Yi. CCS + time = An interleaving model for real time systems. In *ICALP '91*, vol. 510 of *LNCS*, pp. 217–228, 1991.

APPENDIX: PROOF SKETCHES

Given the space constraints, we can only sketch the proofs of our main theorems. The full proofs will be made available in a technical report.

For the proofs of our main theorems of Sec. 4, it will be convenient to first formally specify the class of processes, typified by $Comp$, to which our allowable components will belong.

$$\begin{aligned}
Comp &\stackrel{\text{def}}{=} \{P \mid \forall Q, s \in \mathcal{A}^* \cdot (P \xrightarrow{s} Q) \Rightarrow \\
&\quad \exists R \cdot Q \cong R \in (Inputs(P) \cup Ready(P) \cup Waiting(P) \cup \\
&\quad\quad Internal(P) \cup Outputs(P) \cup Finished(P))\} \\
Inputs(P) &\stackrel{\text{def}}{=} \{Q \mid Q \in Deriv(P) \wedge \exists i \in I \cdot Q \xrightarrow{i} \wedge \\
&\quad \forall \alpha \in \mathcal{A}, R \in \mathcal{P} \cdot Q \xrightarrow{\alpha} R \Rightarrow \alpha \in I \wedge \\
&\quad \exists S \in \mathcal{P} \cdot R \cong S \in (Inputs(P) \cup Ready(P))\} \\
Ready(P) &\stackrel{\text{def}}{=} \{Q \mid Q \in Deriv(P) \wedge Q \xrightarrow{\bar{\tau}} \wedge \\
&\quad \forall \alpha \in \mathcal{A}, R, R' \in \mathcal{P} \cdot (Q \xrightarrow{\alpha} R \wedge Q \xrightarrow{\alpha} R') \Rightarrow \alpha = \bar{\tau} \wedge \\
&\quad \exists S \in \mathcal{P} \cdot R \cong R' \cong S \in Waiting(P)\} \\
Waiting(P) &\stackrel{\text{def}}{=} \{Q \mid Q \in Deriv(P) \wedge Q \xrightarrow{g} \wedge \\
&\quad \forall \alpha \in \mathcal{A}, R, R' \in \mathcal{P} \cdot (Q \xrightarrow{\alpha} R \wedge Q \xrightarrow{\alpha} R') \Rightarrow \alpha = g \wedge \\
&\quad \exists S \in \mathcal{P} \cdot R \cong R' \cong S \in Internal(P)\} \\
Internal(P) &\stackrel{\text{def}}{=} \{Q \mid Q \in Deriv(P) \wedge Q \xrightarrow{\tau} \wedge \\
&\quad \forall \alpha \in \mathcal{A}, R \in \mathcal{P} \cdot (Q \xrightarrow{\alpha} R) \Rightarrow \alpha = \tau \wedge \\
&\quad \exists S \in \mathcal{P} \cdot R \cong S \in Outputs(P)\} \\
Outputs(P) &\stackrel{\text{def}}{=} \{Q \mid Q \in Deriv(P) \wedge \exists o \in O \cdot Q \xrightarrow{o} \wedge \\
&\quad \forall \alpha, \alpha' \in \mathcal{A}, R, R' \in \mathcal{P} \cdot (Q \xrightarrow{\alpha} R \wedge Q \xrightarrow{\alpha'} R') \Rightarrow \alpha = \alpha' \wedge \\
&\quad \exists S \cdot R \cong R' \cong S \in (Outputs(P) \cup Finished(P))\} \\
Finished(P) &\stackrel{\text{def}}{=} \{Q \mid Q \in Deriv(P) \wedge Q \xrightarrow{\bar{g}} \wedge \\
&\quad \forall \alpha \in \mathcal{A}, R, R' \in \mathcal{P} \cdot (Q \xrightarrow{\alpha} R \wedge Q \xrightarrow{\alpha} R') \Rightarrow \alpha = \bar{g} \wedge \\
&\quad \exists S \cdot R \cong R' \cong S \in (Inputs(P) \cup \mathbf{0})\}
\end{aligned}$$

Where we define each of the allowable classes of states using the following relation, intended to find the first representative for each equivalence class:

$$\begin{aligned}
Deriv(P) &\stackrel{\text{def}}{=} \{Q \mid \exists s \in \mathcal{A}^* \cdot P \xrightarrow{s} Q \wedge \\
&\quad \nexists t, u \in \mathcal{A}^+, Q \in \mathcal{P} \cdot t : u = s \wedge P \xrightarrow{t} Q' \wedge Q \cong Q'\}
\end{aligned}$$

We furthermore observe that, due to the statement that all components behaviours are finite state, each of these above sets will be finite for an acceptable description of component behaviour. We also remind ourselves that the

behaviours must be expressed as CCS terms, so we may infer loops of all clocks at every state not a member of $Internal(Comp)$ for a given $Comp$.

Proof Sketch for Theorem 3 (‘Encapsulation Transparency’)

Proof. We first define:

$$\begin{aligned} \mathcal{O} &\stackrel{\text{def}}{=} \mathcal{O}(Comp) \\ \text{Broads} &\stackrel{\text{def}}{=} \prod_{o \in \mathcal{O}} \text{IsoBroad}(o_{\text{inst}}) \\ \text{Wires} &\stackrel{\text{def}}{=} \prod_{o \in \mathcal{O}} \text{IsoWire}(o_{\text{inst}}, \sigma_s, o) \\ P^* &\stackrel{\text{def}}{=} P \mid \Delta_{\{\sigma_s\} \cup \vec{\sigma}} \\ &\vec{\sigma} \text{ as a vector of clocks such that } \forall o \in \mathcal{O}(Comp) \cdot \exists j \cdot \sigma_j = \sigma_{o_{\text{inst}}} \\ &\text{and } \nexists j \cdot \sigma_j \notin \mathcal{O}(Comp) \end{aligned}$$

We allow ourselves to drop implicit parameters, e.g.:

$$\text{InstWrapper}[_] \stackrel{\text{def}}{=} \text{InstWrapper}(CompInst, \text{IsoBroad}, \sigma_s)[Comp]$$

Then one may re-state the theorem as:

$$\text{EncWrapper}[(\text{InstWrapper}[Comp] \mid \text{Wires})/\vec{\sigma}] \cong Comp^*$$

To show this we construct first a (weak) bisimulation:

$\mathcal{R} \stackrel{\text{def}}{=} \mathcal{R}_{Input} \cup \mathcal{R}_{Ready} \cup \mathcal{R}_{Waiting} \cup \mathcal{R}_{Internal} \cup \mathcal{R}_{Output} \cup \mathcal{R}_{Finished}$, where:

$$\begin{aligned} \mathcal{R}_{Input} &\stackrel{\text{def}}{=} \{ \langle \text{EncWrapper}[(\text{InstWrapper}[P] \mid \text{Wires})/\vec{\sigma}], P^* \rangle \mid \\ &\quad P \in \text{Inputs}(Comp) \} \\ \mathcal{R}_{Ready} &\stackrel{\text{def}}{=} \{ \langle \text{EncWrapper}[(\text{InstWrapper}[P] \mid \text{Wires})/\vec{\sigma}], P^* \rangle, \\ &\quad \langle \text{EncWrapper}[\langle (Q \mid r_e \cdot t_c \cdot \bar{g} \cdot \text{Inst} + t_c \cdot \bar{g} \cdot \text{Inst} \mid \\ &\quad \quad \text{Broads}) \setminus \mathcal{O} \cup \{r, g\} \mid \text{Wires})/\vec{\sigma}], P^* \rangle, \\ &\quad \langle \langle (Q \mid t_c \cdot \bar{g} \cdot \text{Inst} \mid \text{Broads}) \setminus \mathcal{O} \cup \{r, g\} \mid \text{Wires})/\vec{\sigma} \mid \\ &\quad \quad \bar{r} \cdot \underline{g} \cdot \bar{t}_c \cdot t_e \cdot \underline{g} \cdot \text{Enc} \setminus \{t_c, r_e\}/\sigma_s, P^* \rangle \mid \\ &\quad \quad P \in \text{Ready}(Comp), P \xrightarrow{\bar{r}} Q \} \\ \mathcal{R}_{Waiting} &\stackrel{\text{def}}{=} \{ \langle \langle (P \mid t_c \cdot \bar{g} \cdot \text{Inst} \mid \text{Broads}) \setminus \mathcal{O} \cup \{r, g\} \mid \text{Wires})/\vec{\sigma} \mid \\ &\quad \underline{g} \cdot \bar{t}_c \cdot t_e \cdot \underline{g} \cdot \text{Enc} \setminus \{t_c, r_e\}/\sigma_s, P^* \rangle \mid P \in \text{Waiting}(Comp) \} \\ \mathcal{R}_{Internal} &\stackrel{\text{def}}{=} \{ \langle \langle (P \mid t_c \cdot \bar{g} \cdot \text{Inst} \mid \text{Broads}) \setminus \mathcal{O} \cup \{r, g\} \mid \text{Wires})/\vec{\sigma} \mid \\ &\quad \bar{t}_c \cdot t_e \cdot \underline{g} \cdot \text{Enc} \setminus \{t_c, r_e\}/\sigma_s, Q^* \rangle, \\ &\quad \langle \langle (P \mid \bar{g} \cdot \text{Inst} \mid \text{Broads}) \setminus \mathcal{O} \cup \{r, g\} \mid \text{Wires})/\vec{\sigma} \mid \\ &\quad \quad t_e \cdot \underline{g} \cdot \text{Enc} \setminus \{t_c, r_e\}/\sigma_s, Q^* \rangle, \\ &\quad \langle \langle (Q \mid \text{Inst} \mid \text{Broads}) \setminus \mathcal{O} \cup \{r, g\} \mid \text{Wires})/\vec{\sigma} \mid \\ &\quad \quad t_e \cdot \underline{g} \cdot \text{Enc} \setminus \{t_c, r_e\}/\sigma_s, Q^* \rangle \mid \\ &\quad \quad P \in \text{Waiting}(Comp), P \xrightarrow{g} Q \} \end{aligned}$$

$$\begin{aligned}
 \mathcal{R}_{Outputs} \stackrel{\text{def}}{=} & \{ \langle (((P \mid \text{Inst} \mid \text{Broads}) \setminus \mathcal{O} \cup \{r, g\} \mid \text{Wires}) / \vec{\sigma}) \\
 & \quad \mid t_e \cdot \underline{g}_{\sigma_s} \cdot \text{Enc}) \setminus \{t_c, r_e\} / \sigma_s, P^* \rangle, \\
 & \langle (((Q \mid \overline{b}_c \cdot f_c \cdot \text{Inst} \mid \text{Broads}) \setminus \mathcal{O} \cup \{r, g\} \mid \text{Wires}) / \vec{\sigma}) \\
 & \quad \mid t_e \cdot \underline{g}_{\sigma_s} \cdot \text{Enc}) \setminus \{t_c, r_e\} / \sigma_s, P^* \rangle, \\
 & \langle (((Q \mid f_c \cdot \text{Inst} \mid \Pi_{o \in \mathcal{O} \setminus c} \text{IsoBroad}(o_{\text{inst}}) \\
 & \quad \mid \text{IsoBroad}'(c_{\text{inst}})) \setminus \mathcal{O} \cup \{r, g\} \\
 & \quad \mid \text{Wires}) / \vec{\sigma}) \mid t_e \cdot \underline{g}_{\sigma_s} \cdot \text{Enc}) \setminus \{t_c, r_e\} / \sigma_s, P^* \rangle, \\
 & \langle (((Q \mid f_c \cdot \text{Inst} \mid \Pi_{o \in \mathcal{O} \setminus c} \text{IsoBroad}(o_{\text{inst}}) \\
 & \quad \mid \text{IsoBroad}'(c_{\text{inst}})) \setminus \mathcal{O} \cup \{r, g\} \\
 & \quad \mid \Pi_{o \in \mathcal{O} \setminus c} \text{IsoWire}(o_{\text{inst}}, \sigma_s, o) \\
 & \quad \mid \underline{t}_{\{\sigma_s, \sigma_{c_{\text{inst}}}\}} \cdot \underline{\sigma}_{c_{\text{inst}}} \cdot \text{IsoWire}(c_{\text{inst}}, \sigma_s, c)) / \vec{\sigma}) \\
 & \quad \mid t_e \cdot \underline{g}_{\sigma_s} \cdot \text{Enc}) \setminus \{t_c, r_e\} / \sigma_s, P^* \rangle, \\
 & \langle (((Q \mid f_c \cdot \text{Inst} \mid \Pi_{o \in \mathcal{O} \setminus c} \text{IsoBroad}(o_{\text{inst}}) \\
 & \quad \mid \text{IsoBroad}'(c_{\text{inst}})) \setminus \mathcal{O} \cup \{r, g\} \\
 & \quad \mid \Pi_{o \in \mathcal{O} \setminus c} \text{IsoWire}(o_{\text{inst}}, \sigma_s, o) \mid \underline{\sigma}_{c_{\text{inst}}} \cdot \text{IsoWire}(c_{\text{inst}}, \sigma_s, c)) / \vec{\sigma}) \\
 & \quad \mid t_e \cdot \underline{g}_{\sigma_s} \cdot \text{Enc}) \setminus \{t_c, r_e\} / \sigma_s, Q^* \rangle, \\
 & \langle (((Q \mid f_c \cdot \text{Inst} \mid \Pi_{o \in \mathcal{O} \setminus c} \text{IsoBroad}(o_{\text{inst}}) \\
 & \quad \mid \overline{f}_c \cdot \text{IsoBroad}(c_{\text{inst}})) \setminus \mathcal{O} \cup \{r, g\} \mid \text{Wires} \\
 & \quad \mid t_e \cdot \underline{g}_{\sigma_s} \cdot \text{Enc}) \setminus \{t_c, r_e\} / \sigma_s, Q^* \rangle \mid \\
 & \quad P \in \text{Output}(\text{Comp}), P \xrightarrow{\vec{g}} Q \} \\
 \mathcal{R}_{Finished} \stackrel{\text{def}}{=} & \{ \langle (((P \mid \text{Inst} \mid \text{Broads}) \setminus \mathcal{O} \cup \{r, g\} \mid \text{Wires} \\
 & \quad \mid t_e \cdot \underline{g}_{\sigma_s} \cdot \text{Enc}) \setminus \{t_c, r_e\} / \sigma_s, P^* \rangle, \\
 & \langle (((Q \mid \overline{t}_c \cdot \text{CompInst} \mid \sigma_s(\overline{t}_c \cdot \text{CompInst}) \\
 & \quad \mid \text{Broads}) \setminus \mathcal{O} \cup \{r, g\} \mid \text{Wires} \\
 & \quad \mid t_e \cdot \underline{g}_{\sigma_s} \cdot \text{Enc}) \setminus \{t_c, r_e\} / \sigma_s, P^* \rangle, \\
 & \langle (((Q \mid \overline{t}_e \cdot \text{CompInst} \mid \text{Broads}) \setminus \mathcal{O} \cup \{r, g\} \mid \text{Wires} \\
 & \quad \mid t_e \cdot \underline{g}_{\sigma_s} \cdot \text{Enc}) \setminus \{t_c, r_e\} / \sigma_s, P^* \rangle, \\
 & \langle (((Q \mid \text{CompInst} \mid \text{Broads}) \setminus \mathcal{O} \cup \{r, g\} \mid \text{Wires} \\
 & \quad \mid \underline{g}_{\sigma_s} \cdot \text{Enc}) \setminus \{t_c, r_e\} / \sigma_s, P^* \rangle, \\
 & \langle (((\mathbf{0} \mid \text{CompInst} \mid \text{Broads}) \setminus \mathcal{O} \cup \{r, g\} \mid \text{Wires} \\
 & \quad \mid \text{Enc}) \setminus \{t_c, r_e\} / \sigma_s, \mathbf{0}^* \rangle \mid \\
 & \quad P \in \text{Finished}(\text{Comp}), P \xrightarrow{\vec{g}} Q \}
 \end{aligned}$$

From this bisimulation we show congruence since all clocks but those in the set $\{\sigma_s\} \cup \vec{\sigma}$ idle on both sides of the equivalence until an input is received at which point it is only necessary, from the definition of the congruence, to show observation equivalence.

Proof Sketch for Theorem 2 ('Isochronic Forks')

Proof. We first construct a weak bisimulation as follows:

$$\{ \langle (\text{IsoBroad}(o) \mid \Pi_{n \leq |\vec{i}|} \text{IsoWire}(o, \sigma_s, i_n)) / \sigma_o \mid \Delta_{\sigma_s}, \text{Fork}(o, \vec{i}) \mid \Delta_{\{\sigma_s, \sigma_o\}} \rangle \}$$

U

$$\begin{aligned} & \{ \langle \text{IsoBroad}'(o) \mid \Pi_{n \leq |\vec{j}|} \underline{\sigma}_o \cdot \text{IsoWire}(o, \sigma_s, j_n) \\ & \quad \mid \Pi_{n \leq |\vec{k}|} \overline{k_n}_{\{\sigma_s, \sigma_o\}} \cdot \underline{\sigma}_o \cdot \text{IsoWire}(o, \sigma_s, k_n) \\ & \quad \mid \Pi_{n \leq |\vec{l}|} \text{IsoWire}(o, \sigma_s, l_n) \rangle / \sigma_o \mid \Delta_{\sigma_s}, \text{Fork}'(o, \vec{i}, (\vec{i} \setminus \vec{j})) \mid \Delta_{\{\sigma_s, \sigma_o\}} \rangle \mid \\ & \quad \mid \vec{j} \mid + \mid \vec{k} \mid + \mid \vec{l} \mid = \mid \vec{i} \mid \wedge \forall n \leq \mid \vec{i} \mid \cdot (\exists m \cdot j_m = i_n) \vee (\exists m \cdot k_m = i_n) \vee (\exists m \cdot l_m = i_n) \} \\ & \text{(Abusing the notation to form } \vec{i} \setminus \vec{j} \text{ meaning all those members of } \vec{i} \text{ not already} \\ & \text{included in } \vec{j}, \text{ i.e., not having been already communicated.)} \end{aligned}$$

U

$$\{ \langle \overline{f}_o \cdot \text{IsoBroad}(o) \mid \Pi_{n \leq |\vec{i}|} \text{IsoWire}(o, \sigma_s, i_n) \rangle / \sigma_o \mid \Delta_{\sigma_s}, \overline{f}_o \cdot \text{Fork}(o, \vec{i}) \mid \Delta_{\{\sigma_s, \sigma_o\}} \rangle \}$$

We then observe that both agents idle on all clocks except σ_s and σ_o until the a -transition is taken and are therefore congruent.

Proof Sketch for Theorem 4 ('Jam Check')

Proof. From the definition of Rule (tTO1), we see that $\text{Check} \stackrel{\text{def}}{=} \mu x. [\Delta]_{\sigma_s}(x)$ defines a state with a self-transition in σ_s only. Under the assumptions of the theorem, System produces transitions in labels σ_s and τ only. Since all σ_s transitions of Check may be matched weakly under observation equivalence, any system in which σ_s is live can therefore be shown equivalent. On the other hand, a system where there exists a path to a state where σ_s is indefinitely stalled clearly cannot match successive ticks of σ_s and therefore will not be equivalent.