

TUM

INSTITUT FÜR INFORMATIK

Compatibility and reuse in component-based
systems via type and unit inference

Christian Kühnel, Andreas Bauer, Michael Tautschnig



TUM-I0716

Mai 07

TECHNISCHE UNIVERSITÄT MÜNCHEN

TUM-INFO-05-I0716-0/1.-FI

Alle Rechte vorbehalten

Nachdruck auch auszugsweise verboten

©2007

Druck: Institut für Informatik der
 Technischen Universität München

Compatibility and reuse in component-based systems via type and unit inference*

Christian Kühnel¹

Andreas Bauer^{2,1}

Michael Tautschnig¹

¹Institut für Informatik, Technische Universität München

²National ICT Australia (NICTA), Canberra

Abstract

In many branches of industry, the component-based approach to systems design is predominant, e. g., as in embedded control systems which are often modelled using MATLAB/Simulink. In order to facilitate reuse, and to raise the level of abstraction for future designs and frequently used functions, the employed tool sets offer built-in mechanisms to create sophisticated component libraries. For large, real-world designs, however, it is not always clear, whether or not a certain context violates even the most basic design assumptions of employed library components, thus often leading to expensive runtime errors. This paper introduces a practical method for checking compatibility of large designs, statically. This method not only ensures that large component-based designs provide a context such that all (library) components have well defined types, but it also ensures that transmitted physical units, such as m^2 , km/h , mph , etc. are preserved during computation. As such the possibility for runtime errors is reduced, and a metric for sound component reuse given.

1 Introduction

Component-based approaches to software and systems design are predominant in large parts of industry, such as automotive and embedded control systems in general. Tool chains used in these domains are often based on MATLAB/Simulink, a graphical modelling and simulation environment with the ability to generate code (via program extensions). MATLAB/Simulink and other tools further support the creation and use of dedicated component libraries such that frequently used functionality can be reused in different contexts, and without having to concentrate on implementation details. As such, the reuse of components raises

the level of abstraction for many designs, and development time is often reduced.

However, this convenience comes at a price; that is, the developer has to make sure that reused components are compatible in the context they are used in. For large, real-world designs as they are common, e. g., for automotive control systems like adaptive cruise control, engine management, or electronic brake systems, reuse is often constrained by the ability to establish firm criteria for ensuring compatibility of reused components with respect to the rest of a complex design, i. e., a *context*.

Contribution. In this paper, we focus on establishing a criterion for *syntactic compatibility*, i. e., a criterion that can be checked at compile-time and without executing the component-based design under scrutiny. In the context of some of our industrial research collaborations, compatibility checking has been reduced to a problem of *type inference* for which a comprehensive theory already exists (cf. [13]), and which can be greatly automated and tool-supported. We say that a given design is compatible w. r. t. its employed library components iff (1.) a well-typed design can be inferred from it, and (2.) the inferred types do not violate a set of predetermined *measurement units*. Measurement units (see Sec. 4) play an important role in domains where systems control physical processes, e. g., for ensuring that signals are interpreted correctly by another system or component. Therefore, a *decidable polymorphic* type system for a custom component language is introduced that can be used to infer for a concrete design its types and according measurement units, and determine compatibility. However, this type system and inference algorithm is not restricted to our language, but can be applied to similar component-based formalisms, e. g., MATLAB/Simulink at ease.

In the context of our current projects, our approach to compatibility testing has proven to be a reliable metric not only in terms of establishing type compatibility of large designs and component libraries, but also as a means of avoiding type errors in the running program and the code gener-

*This report is an extended version of [10].

ator of our component language. An implementation of our method would exhibit only linear runtime in the number of design artifacts due to the use of standard *unification* algorithms (cf. [9]) for inference. Moreover, not only type correctness can be established formally, but also compatibility between the measurement units of a design.

Related work. Most tools which support a component or model-based design approach, and which are currently used in the domain of embedded control systems, do not offer advanced concepts for type checking, let alone type inference. The authors of this text believe that the reason lies in that these systems target rather low-level design artifacts, such as the design of continuous control algorithms, and abstraction from this level of detail has not been the primary focus of the developers for a long time.

The concept of measurement units in combination with a strong static type system, however, caters for comprehensible designs, and avoids errors in the actually running systems. Moreover, they facilitate the reuse of components, and give to the user a powerful tool for combining different library components in different contexts without violating assumptions over measurement units or data types, and not being aware of it.

Currently, MATLAB/Simulink does not offer strong static type checking, and subsumes mostly basic data types such as `boolean` and `double` which, if not specified at “compile-time”, may create runtime errors due to multyped designs; that is, if nothing is explicitly specified for a component, then `double` is used as a “default type” for input and output signals (cf. [5, 1]). The situation is a similar one for other domain-specific solutions such as ASCET-SD.

However, various authors have concentrated on establishing *behavioural correctness* of low-level designs. A popular formalism are *Interface automata* [6]. These can be used for the specification of behaviour at component interfaces, and to establish compatibility between compound interfaces by means of automata analysis. One of the research tools already making use of this is, e.g., Ptolemy II [11]. The static part of Ptolemy’s type system, however, is rather different from the approach presented in this paper, in that the behaviour of Ptolemy II components (called “actors”) is tightly interwoven with Java code. Therefore, type correctness is partly handled by the Java compiler, and Java is known to leave space for subtle type errors at runtime.

Our custom language, in the next section introduced as SCL, handles type inference at compile time and avoids such problems altogether. To the best of the knowledge of the authors, no other component-based design language for the development of embedded control systems currently offers the combination of polymorphic type and measurement unit inference as well as a decidable, static type system.

2 SCL—A simple component language

In the following, we introduce a simple component language, SCL, that provides the essential mechanisms for systems modelling and analysis. Conceptually, SCL is close to other graphical modelling formalisms, such as UML-based ones, or MATLAB/Simulink, but its semantics is based upon a uniform discrete time-base and the hypothesis of perfect-synchrony [2]. This basically asserts that computations of components occur instantly, i.e., take no time, and that communication is infinitely fast. This allows us to abstract from implementation details when modelling, and caters for a sound and well-defined semantics. Various tools for systems design and synthesis exist that are based on this notion, e.g., SCADE [5], Esterel [3], and AutoFocus [4].

In the following the syntax and semantics of SCL are introduced. The syntax is required for the definition of the typing rules (see Sec. 3.1), the semantics for the safety proof (see Sec. 3.2).

2.1 Concepts and syntax

Basically, SCL provides the following core concepts for modelling: (functional) blocks, ports and channels for communication, (composite) components, all of which are explained in the following.




Figure 1. An example SCL model

A block is a “primitive” component whose functionality is defined by the semantics of SCL. It cannot be decomposed further. For instance, the predefined arithmetic block *plus* simply adds two variables, which are delivered to it via two input ports (see Fig. 1). Further, a dedicated *delay* block (i.e., z^{-1}) is used to store values for one computational cycle, and essential for avoiding causal loops in the model. Hence its value gets updated periodically, and needs to be initialised to a default value when invoked for the first time. For brevity, only some predefined blocks are discussed here.

A block has a signature S and parameters depending on the type of block, e.g., the *delay* block $delay((p_i \mapsto p_o), v)$ has an initial value v . (Composite) components, denoted $component((S), ch, C, \Phi)$, are composed of a nonempty set of subcomponents C , and a nonempty set

of channels ch , and have a well-defined signature $S = p_1^{in}, \dots, p_n^{in} \mapsto p_1^{out}, \dots, p_n^{out}$ that is made up of its outside-visible input and output ports. The type context Φ is used to introduce new type variables.

The syntax of SCL is defined in Fig. 2 and Fig. 3.

p	port
v	value
T	type
Γ	type context

Figure 2. Syntax of SCL– primitive elements

2.2 Operational semantics

In order to substantiate the claims made in Sec. 1, we briefly introduce an operational semantics for SCL. This provides for a concrete type context when reasoning about type safety, and consequently compatibility of component-based designs in SCL.

The notation for the operational semantics used here is loosely based on the one presented by Plotkin [14]. The rules are displayed in Fig. 4. The current state of a system is denoted by σ , and the history by η (i. e., the previous state). This differentiation is important especially in face of the *delay* block which stores values for exactly one computational cycle. We fix a set of ports P , and a set of channels ch , as well as a domain of values V which are transmitted via ports and channels. Let $p_1, p_2 \in P$ and $c \in ch$, then $c = chan(p_1 \mapsto p_2)$ indicates that a port p_1 is connected with p_2 via channel c . Further, we use the notation $\sigma(p) \in V$ for “retrieving” a value currently associated to a port $p \in P$, and $\sigma[v/p]$ for “storing” it in the state (respectively with η). To avoid conflicts in σ and η , the identifiers of all ports and delays of a model are assumed to be unique.

Finally, we introduce an artificial function $ord : 2^E \rightarrow 2^E$, which, for a given set of entities, returns a partially ordered set of entities that defines their order of execution, where $E = C \cup ch$ and C is the set of components connected over the set of channels ch . Before a component can be evaluated, the input ports have to be assigned their values. For the system, this has to be provided by the environment, within a component this is ensured by the *ord* function.

In the first computational cycle (i. e., at the beginning of execution), the history is empty, $\eta = \emptyset$. In the following cycles, the history of step n is used to compute step $n + 1$.

Order of execution and causality The values of a component (S-comp) are defined according to its subcomponents and channels. The execution order of the individual subcomponents and channels is determined by function *ord*.

Additionally, the *ord*-function formally “splits” every *delay* (S-delay) into a load part, $delay_l$, and a store part, $delay_s$. The load part is placed at the beginning of the evaluation order, and the store part at the end. This ensures that the delayed values of the last step are available to all components and that the current values of the delays are stored.

The remaining components and the channels are sorted in their evaluation order. Let $c \in C$ be some component, and let $I(c)$ denote the set of input ports of c and $O(c)$ the output ports, respectively, then a partial order $<$ is defined as follows: For two components $c, c' \in C$, and for every channel $chan(p_o \mapsto p_i) \in ch$, $c < chan(p_o \mapsto p_i) < c'$ iff $p_o \in O(c)$ and $p_i \in I(c')$. Since no cyclic connections without a *delay* are allowed, $<$ delivers a partial order on model entities, returned by *ord*.

Let $;$ be the functional sequence operator. Then, an ordered sequence $e; es$ of entities is evaluated by evaluating the first element e of the sequence, delivering a new valuation sequence σ' and history sequence η' . These new sequences are then used for recursive evaluation of the remainder of the entity sequence es as depicted in rule (S-seq).

The store part of a delay takes the value $v \in V$ of the input port $p_i \in P$ and stores it in the history η for the next cycle (S-delay-store). In the first cycle, when $\eta(p) = \perp$, i. e., the port is not yet bound in η , the load part of delay block returns its default value (S-delay-nohist). In the consecutive steps, it returns a value $v_o \in V$ stored η (S-delay-hist).

As expected, a channel simply propagates the value of one port to the other port (S-chan).

Execution of blocks. Let in what follows, $p_a, p_b, p_c \in P$ and $v \in V$. A *plus* block is evaluated in (S-plus) by adding the values of the input ports p_a and p_b and binding the result to the output port p_c . As we are using existential types for type abstraction, (see Sec. 3.1), the *plus*-Operator of the respective type is applied. The rules for other blocks such as *minus*, *mult*, *div*, *and*, *or*, etc. are similar. An entry l of an existentially typed value v can be accessed by $v.l$ (S-rcd).

The evaluation of a conditional block, *if*, depends on the value of port p_c : if $\sigma(p_c) = true$ then the value of the port p_t is assigned to the result port p_r (S-if-true), otherwise the value of port p_e is used (S-if-false).

3 Compatibility in SCL designs

Compatibility checking in SCL is reduced to three key aspects all of which can be checked statically and automatically:

1. Type inference, i. e., inferring the appropriate types for an under-specified model, if syntactically possible.

$S :=$	$p, \dots, p \rightsquigarrow p, \dots, p$	signature
$ch :=$	$chan(p, p)$	channel
$E :=$	C	entities
	ch	component
		channel
$C :=$	$component((S), 2^C, 2^{ch}, I, \Gamma)$	components
	B	hierarchical component
		block
$B :=$	$delay((S), v)$	blocks
	$const(p, v)$	$delay$ block
	$plus(S)$	constant
	$if(S)$	$plus$ block
	\dots	if block

Figure 3. Syntax of SCL– composed elements

$\frac{\langle ord(ch \cup C), \sigma, \eta \rangle \rightarrow \langle \sigma', \eta' \rangle}{\langle component(S, ch, C, \Phi), \sigma, \eta \rangle \rightarrow \langle \sigma', \eta' \rangle}$	(S-comp)	$\frac{\eta(p_i) = \perp}{\langle delay_t((p_i \rightsquigarrow p_o), v), \sigma, \eta \rangle \rightarrow \langle \sigma[v/p_o], \eta \rangle}$	(S-delay-nohist)
$\frac{\langle delay_t((p_i \rightsquigarrow p_o), v), \sigma, \eta \rangle \rightarrow \langle \sigma', \eta' \rangle \quad \langle delay_s((p_i \rightsquigarrow p_o), v), \sigma, \eta \rangle \rightarrow \langle \sigma, \eta' \rangle}{\langle delay((p_i \rightsquigarrow p_o), v), \sigma, \eta \rangle \rightarrow \langle \sigma', \eta' \rangle}$	(S-delay)	$\frac{\eta(p_i) = v_o}{\langle delay_t((p_i \rightsquigarrow p_o), v), \sigma, \eta \rangle \rightarrow \langle \sigma[v_o/p_o]; \eta \rangle}$	(S-delay-hist)
$\frac{\langle e, \sigma, \eta \rangle \rightarrow \langle \sigma', \eta' \rangle \quad \langle es, \sigma', \eta' \rangle \rightarrow \langle \sigma'', \eta'' \rangle}{\langle e; es, \sigma, \eta \rangle \rightarrow \langle \sigma'', \eta'' \rangle}$	(S-seq)	$\frac{\sigma(p_1) = v}{\langle chan(p_1 \rightsquigarrow p_2), \sigma, \eta \rangle \rightarrow \langle \sigma[v/p_2], \eta \rangle}$	(S-chan)
$\frac{\sigma(p_a) = v_a \quad \sigma(p_b) = v_b \quad v_c.value = v_a.plus(v_a.value, v_b.value)}{\langle plus(p_a, p_b \rightsquigarrow p_c), \sigma, \eta \rangle \rightarrow \langle \sigma[v_c/p_c], \eta \rangle}$	(S-plus)	$\frac{\sigma(p_c) = true \quad \sigma(p_t) = v}{\langle if(p_c, p_t, p_e \rightsquigarrow p_r), \sigma, \eta \rangle \rightarrow \langle \sigma[v/p_r], \eta \rangle}$	(S-if-true)
$\frac{\sigma(p_i) = v_i}{\langle delay_s((p_i \rightsquigarrow p_o), v), \sigma, \eta \rangle \rightarrow \langle \sigma, \eta[v_i/p_i] \rangle}$	(S-delay-store)	$\frac{\sigma(p_c) = false \quad \sigma(p_e) = v}{\langle if(p_c, p_t, p_e \rightsquigarrow p_r), \sigma, \eta \rangle \rightarrow \langle \sigma[v/p_r], \eta \rangle}$	(S-if-false)
		$\frac{v = \{l_1 = v_1, \dots, l_n = v_n\} \quad l_i \in \{l_1, \dots, l_n\}}{v.l_i = v_i}$	(S-rcd)

Figure 4. Operational semantics of SCL

2. Type safety, i. e., SCL ensures that a well-typed model cannot create runtime errors which are due to type failures.
3. Units, i. e., SCL models can be checked for compatibility in terms of the physical units that are computed and used in a model.

In this section, we first concentrate on 1. and 2. The introduction and elimination rules and some of the subtyping rules are omitted here for the sake of brevity, as those are similar to the standard literature (cf. [13]).

3.1 Inference of data types and signatures

When components are built for reuse, it is desirable to design them so that they can be used in a variety of contexts. For example, a sorting algorithm could be used with any type, as long as an ordering, e. g., an \leq -operator, is defined on that type. Another example would be a component implementing a queue: this component should be able to queue elements of a static, but arbitrary type. To implement this in the type system, *existential types*, *universal types* and *subtyping* are used.

If type S is a subtype of type T , denoted $S \sqsubseteq T$, this means that every type T in a well-typed model could be replaced with an S , and the model would still be well-typed.

(Bounded) universal types are used to introduce new type variables in the context Φ of a component (T-component). They are denoted by $\forall X \sqsubseteq T. X$ to express that the component can be used for any type X , which is required to be a subtype of T . This can be used to define the types of the queue example above.

Existential types are used to denote which operations have to be defined on a certain type. In combination with universal types and subtyping, the example can now be described formally: the sort component may be used with any type $\forall X \sqsubseteq T. X$ as long as an operation \leq is defined on that type $T = \{\exists Y, \{\leq: Y \times Y \rightarrow \text{Bool}\}\}$.

When a new type such as a record or tuple is defined by the users, they can also define operations on that type and thereby build a new existential type. This new type can now be used with any (prior) component whose type restrictions it satisfies.

SCL requires only a minimal set of primitive types, `Bool` and `Float`. `Bool` is required for comparison and *if* blocks, `Float` for arithmetics. Other types can be added in a similar manner. As we use only existential types throughout the

model, these types are defined as:

$$\begin{aligned} \text{Float} &:= \{\exists X, \{value : X, \\ &\quad plus, minus, mult, div : X, X \rightarrow X, \\ &\quad sqrt : X \rightarrow X, \\ &\quad eq, le, gt : X, X \rightarrow \text{Bool}\}\} \\ \text{Bool} &:= \{\exists X, \{value : X, \\ &\quad and, or : X, X \rightarrow X, \\ &\quad neg : X \rightarrow X\}\} \end{aligned}$$

Notably, SCL supports under-specification in models in terms of polymorphic blocks and components. For instance, a *plus* block is polymorphic in that it works over different value domains, and its output type is determined solely by the types of input values. A model can thus be made up only of polymorphic blocks (i. e., under-specified blocks and components lacking a concrete signature), and the concrete instances are then inserted automatically by the type inference mechanism.

The (almost) complete set of typing rules of SCL is depicted in Fig. 5. A type T of a port $p \in P$ can be derived from a *type context* Γ , whenever there is exactly one occurrence of p in Γ (T-port). The set of all bound ports in Γ is defined as $dom(\Gamma)$. And whenever a typing can be derived from a context Γ' , so can it from any permutation thereof (T-perm).

The ports of the *plus* block are bound to the type T , which may be an arbitrary type for which an operator *plus* with signature $X, X \rightarrow X$ is defined (T-plus). The rules for *minus*, *mult*, *neg*, and *sqrt*, etc. are similar.

For a component to be well-typed, all of its channels and sub-components have to be well-typed (T-component). A component is the only possibility for the user to define new type variables by adding them to the additional context Φ . The scope of these variables is the component and all its ports and sub-components.

A channel is the “link” between two components. It is well-typed, if the type of the source port $p_1 \in P$ is a subtype of the type of the target port $p_2 \in P$ as defined in rule (T-chan). Naturally a channel can be used for any type satisfying the subtype property. It is sufficient to allow subtyping only at channels, i. e., between components, as the type system would not be more expressive if we allowed subtyping within blocks.

Subtyping. The types in SCL together with the subtyping relation \sqsubseteq form a lattice. The subtyping rules over this lattice are depicted in Fig. 5.

The subtyping relation is transitive (U-trans) and reflexive (U-refl) and every type is a subtype of `Any` (U-any). An existential type is subtype of another existential type, if their operations are in a subtype relation (U-ex). The

$\frac{p \notin \text{dom}(\Gamma)}{\Gamma, p : T \vdash p : T}$	(T-port)	$\frac{\Gamma \vdash p_a, p_b, v : T}{\Gamma \vdash \text{delay}((p_a \rightarrow p_b), v)}$	(T-delay)
$\frac{\Gamma' \vdash p : T \quad \Gamma \text{ permutation of } \Gamma'}{\Gamma \vdash p : T}$	(T-perm)	$\frac{\forall c \in C. \Gamma, \Phi \vdash c \quad \forall c \in \text{ch}. \Gamma, \Phi \vdash c}{\Gamma \vdash \text{component}(S, \text{ch}, C, \Phi)}$	(T-component)
$\frac{\Gamma \vdash p_a, p_b, p_c : T \quad \Gamma \vdash T \sqsubseteq \{\exists X, \{value : X, plus : X, X \rightarrow X\}\}}{\Gamma \vdash plus(p_a, p_b \rightsquigarrow p_c)}$	(T-plus)	$\frac{\Gamma \vdash p_1 : S \quad \Gamma \vdash p_2 : S \quad \Gamma \vdash S \sqsubseteq T}{\Gamma \vdash \text{chan}(p_1 \rightsquigarrow p_2)}$	(T-chan)
$\frac{\Gamma \vdash p_a, p_b, p_c : T \quad \Gamma \vdash T \sqsubseteq \{\exists X, \{value : X, minus : X, X \rightarrow X\}\}}{\Gamma \vdash minus(p_a, p_b \rightsquigarrow p_c)}$	(T-minus)	$\frac{\Gamma \vdash p_c : \text{Bool} \quad \Gamma \vdash p_t, p_e, p_r : T}{\Gamma \vdash \text{if}(p_c, p_t, p_e \rightsquigarrow p_r)}$	(T-if)
		$\overline{\Gamma \vdash T \sqsubseteq T}$	(U-refl)
		$\frac{\Gamma \vdash S \sqsubseteq T \quad \Gamma \vdash T \sqsubseteq Q}{\Gamma \vdash S \sqsubseteq Q}$	(U-trans)
$\frac{\Gamma \vdash p_a, p_b : T \quad \Gamma \vdash p_c : \text{Bool} \quad \Gamma \vdash T \sqsubseteq \{\exists X \{value : X, eq : X, X \rightarrow \text{Bool}\}\}}{\Gamma \vdash eq(p_a, p_b \rightsquigarrow p_c)}$	(T-eq)	$\overline{\Gamma \vdash T \sqsubseteq \text{Any}}$	(U-any)
$\frac{\Gamma \vdash p_a, p_b : T \quad \Gamma \vdash p_c : \text{Bool} \quad \Gamma \vdash T \sqsubseteq \{\exists X \{value : X, gt : X, X \rightarrow \text{Bool}\}\}}{\Gamma \vdash gt(p_a, p_b \rightsquigarrow p_c)}$	(T-gt)	$\frac{\Gamma \vdash S \sqsubseteq R}{\Gamma \vdash \{\exists X, S\} \sqsubseteq \{\exists X, R\}}$	(U-ex)
		$\frac{\{k_j^{j \in 1..m}\} \supseteq \{l_i^{i \in 1..n}\} \quad k_j = l_i \rightarrow S_j \sqsubseteq T_j}{\{k_j : S_j^{j \in 1..m}\} \sqsubseteq \{l_i : T_i^{i \in 1..n}\}}$	(U-rcd)

Figure 5. Typing (T-) and subtyping (U-) rules of SCL

subtyping-relation for records (U-rcd) is used for existential types, since we want

$$\{\exists X, \{value : X, plus, minus : X, X \rightarrow X\}\} \\ \sqsubseteq \{\exists X, \{value : X, plus : X, X \rightarrow X\}\}$$

This case occurs, e. g., when one output port is connected to the input ports of a *plus* and a *minus* block by two channels. Then the type of that output port has to provide a *plus* and a *minus* operator, thus the two existential types have to be unified as defined in rules (U-ex) and (U-rcd).

These properties of the type system not only provide the possibility to check the compatibility of a component but can also be used to adjust the typing of the components to make them compatible. This is achieved by inferring a valid typing from the composed components, if such a typing is possible.

3.2 Runtime safety

Yet, establishing the compatibility of two components at compile time is not enough. The main concern is that these components remain compatible at runtime. This means that the type system must ensure that no type errors occur at runtime. We therefore have to prove *safety* of our static type system. This means that a model, which was well-typed at compile-time, remains well-typed during runtime.

In general, to show safety for a type system it is sufficient to show *progress* and *preservation* [13]. Progress means that the evaluation, based on the operational semantics, for any well-typed entity is not stuck, i. e., a port can be evaluated and for all entities, there is a rule that can be applied. Preservation means that every well-typed entity remains well-typed after an evaluation step in the operational semantics. Syntactic correctness of the model is presumed here.

Lemma 1 (Progress) *If the model and σ and η are well-typed, then every port can be evaluated and for each entity there is a rule of the operational semantics that can be applied to it.*

Before we proceed to show safety, we introduce two lemmas which follow as a direct consequence of the syntax and semantics of SCL.

Lemma 2 *For any syntactically correct model, all of the input ports of every entity e have been assigned values in σ before e is evaluated.*

Lemma 3 *In every tick but the first, the input port p_i of every delay block is bound in η .*

Safety can be shown in terms of two properties of the type system: progress and preservation. In the following, we show both.

3.3 Proof for Lemma 1 (Progress)

We show this by induction over the rules of the operational semantics as shown in Fig. 4.

case (S-plus) By Lemma 2, p_a and p_b are bound in σ and therefore $\sigma(p_a) = v_a$ and $\sigma(p_b) = v_b$ can be evaluated. As the block and σ are well typed, the ports p_a, p_b, p_c and their values v_a, v_b have type T , with $T \sqsubseteq \{\exists X, \{value : X, plus : X, X \rightarrow X\}\}$. Thus v_a is an existential type, with a defined *plus* operator and both v_a and v_b have *value* fields. Therefore $v_c.value = v_a.plus(v_a.value, v_b.value)$ can be evaluated and v_c can be assigned to p_c in σ . Thus $\langle plus(p_a, p_b \mapsto p_c), \sigma, \eta \rangle$ can be evaluated to $\langle \sigma[v_c/p_c], \eta \rangle$.

case (S-chan) By Lemma 2, p_a is bound in σ and therefore $\sigma(p_1) = v$ can be evaluated. As $p_1, p_2 : T$ and σ are well typed, so is $v : T$. Thus v can be assigned to p_2 in σ and with that $\langle chan(p_1 \mapsto p_2), \sigma, \eta \rangle$ can be evaluated to $\langle \sigma[v/p_2], \eta \rangle$.

case (S-delay-hist), (S-delay-nohist) For the load-part of the delay two cases have to be distinguished: in the first tick, the history η is empty, in the following ticks it is nonempty, according to Lemma 3.

subcase Thus in the first tick (S-delay-nohist) can be applied as $p_o, v : T$ are well typed and v can be assigned to p_o in σ .

subcase In all following ticks, rule (S-delay-hist) is used. As η is well-typed, p_i, v_o are of type T . Thus v_o can be assigned to p_o in σ .

Thus $\langle delay_l((p_i \mapsto p_o), v), \sigma, \eta \rangle$ can be evaluated to $\langle \sigma[v/p_o], \eta \rangle$ or $\langle \sigma[v_o/p_o], \eta \rangle$ respectively.

case (S-delay-store) By Lemma 2, p_i is bound in σ and therefore $\sigma(p_i)$ can be evaluated. As $p_i, v_i : T$ are well typed, v_i can be assigned to p_i in η . Thus $\langle delay_s((p_i \mapsto p_o), v), \sigma, \eta \rangle$ can be evaluated to $\langle \sigma; \eta[v_i/p_i] \rangle$

case (S-if-true) and (S-if-false) By Lemma 2, p_c, p_t, p_e are bound in σ and thus $\sigma(p_c), \sigma(p_t), \sigma(p_e)$ can be evaluated. As $p_c : Bool$, its value either is *true* or *false*. If it is *true*, rule (S-if-true) is applied if *false* rule (S-if-false). As $p_t, p_e, p_r : T$ are well typed, the values of either p_t or p_e can be assigned to p_r in σ . Thus $\langle if(p_c, p_t, p_e \mapsto p_r), \sigma, \eta \rangle$ can be evaluated to $\langle \sigma[v/p_r], \eta \rangle$.

case (S-seq) To evaluate $\langle e, \sigma, \eta \rangle$ a matching rule for the entity e has to be found. An entity is either a block or a channel. If it is a channel, rule (S-chan), if it is a block, the respective block-rule can be applied

as shown above. Thus $\langle e, \sigma, \eta \rangle$ can be evaluated to $\langle \sigma', \eta' \rangle$. For es two cases have to be investigated:

subcase If es is an entity sequence, the rules (S-seq) can be applied recursively.

subcase If es is a single entity, one of the entity rules can be applied as shown for e and with that $\langle es, \sigma', \eta' \rangle$ can be evaluated to $\langle \sigma'', \eta'' \rangle$.

Thus $\langle e; es, \sigma, \eta \rangle$ can be evaluated to $\langle \sigma'', \eta'' \rangle$

case (S-comp) As the model is assumed to be syntactically correct, it does not contain any cycles without delays and thus the ord function can compute an execution sequence for all components C and channels ch . This execution sequence $\langle ord(ch \cup C), \sigma, \eta \rangle$ can be evaluated to $\langle \sigma', \eta' \rangle$ using rule (S-seq) as shown above. Thus $\langle component(S, ch, C, \Phi), \sigma, \eta \rangle$ can be evaluated to $\langle \sigma', \eta' \rangle$.

others The cases omitted for brevity can be argued in the same way as the ones presented above. \square

Lemma 4 (Preservation) *If the model is well-typed and for an evaluation rule $\langle e, \sigma, \eta \rangle \rightarrow \langle \sigma', \eta' \rangle$ σ and η are well-typed so are σ' and η' . A model is well-typed, if the typing rules (see Sec. 3.1) can be applied without getting stuck. σ (and η respectively) is well-typed, iff for every binding $\sigma(p) = v$, v and p are of the same type, i. e., $p, v : T$.*

3.4 Proof for Lemma 4 (Preservation)

Again, we show this by induction over the rules of the operational semantics. In each case we assume that σ and η are well-typed and show that σ' and η' are also well-typed. If the binding of one port is changed, the binding of the remaining ports is unchanged, and therefore these remain well-typed.

case (S-plus) As p_a, p_b, p_c and v_a, v_b are well-typed, they of type T , with $T \sqsubseteq \{\exists X, \{value : X, plus : X, X \rightarrow X\}\}$. From that follows that $v_a.value, v_b.value : X$ and $v_a.plus : X, X \rightarrow X$. Thus $v_c.value = v_a.plus(v_a.value, v_b.value)$ is of type X and v_c is also of type T . Thus $\langle \sigma[v_c/p_c], \eta \rangle$ is still well typed.

case (S-chan) If $p_1 : S$ and $p_2 : T$ then $v = \sigma(p_1) : S$ with $S \sqsubseteq T$. So v can be assigned to p_2 and thus $\langle \sigma[v/p_2], \eta \rangle$ is also well-typed.

case (S-delay-nohist) As the model is well-typed p_o and v are of type T and therefore $\langle \sigma[v/p_o], \eta \rangle$ again is well-typed.

case (S-delay-hist) As p_i, p_o and $v_o = \eta(p_i)$ are of type T and thus v_o can be assigned to p_i in σ and thus $\langle \sigma[v_o/p_o], \eta \rangle$ is still well-typed.

case (S-delay-store) As p_i and v_i are of type T and v_i can be assigned to p_i in η . Thus $\langle \sigma, \eta[v_i/p_i], \eta \rangle$ is also well-typed.

case (S-if-true) and (S-if-false) As σ is well-typed, p_t, p_e and p_r are of type T , as well as $v = \sigma(p_t)$ (and $v = \sigma(p_e)$ respectively.) Thus $\langle \sigma[v/p_r], \eta \rangle$ again is well-typed.

case (S-seq) If $\langle \sigma, \eta \rangle$ are well-typed, every entity $\langle e, \sigma, \eta \rangle$ be evaluated to $\langle \sigma', \eta' \rangle$ and $\langle \sigma', \eta' \rangle$ are still well typed as shown above. If es is a entity sequence, the rule (S-seq) is applied recursively, if it is a single entity, the same argumentation as for e can be used to show that $\langle \sigma'', \eta'' \rangle$ is also well-typed.

case (S-comp) The ord function only sorts the entities of a component and does therefore not influence their typing. Thus the resulting entity sequence $\langle ord(ch \cup C), \sigma, \eta \rangle$ is evaluated to a well-typed $\langle \sigma', \eta' \rangle$ as shown for rule (S-seq). Thus is rule also is type preserving.

others The cases omitted for brevity can be argued in the same way as the ones presented above. \square

Theorem 1 (Safety) *This type system is safe, i. e., there are no type errors at run-time.*

Proof Safety can be shown by proving progress and preservation [13]. \square

Thus, when two components are compatible at compile time, they remain compatible during run-time. This is especially important for embedded systems, as a type error at runtime might lead to a damage of the system or its environment.

4 Measurement units

When a port in a model is assigned type Float, all we know is the range of values for that port. It does not say anything about what these values actually represent in the real world. A model in which two integers, say, one representing apples and the other oranges, are added would still be well typed. But in many cases it is not desirable that apples and oranges can be added. In such a case, additional semantic information must be represented in the model. This is achieved using *measurement units*.

Let \mathbb{U} denote the set of all (measurement) units defined for a model. Which units are used within a model

depends on the application domain. For embedded systems the SI units [12] could be used, as many physical values are processed in such an environment, i. e., $\mathbb{U} = \{s, m, kg, A, K, mol, cd\}$. The function $\mu : \mathbb{U} \rightarrow \mathbb{Z}$ is a mapping from units \mathbb{U} to integral numbers \mathbb{Z} representing the exponent of the respective unit. This is necessary as mixed units, such as $acceleration = m/s^2$, can occur as well. This example would be represented as $\mu(m) = 1$, $\mu(s) = -2$ and $\mu(u) = 0$ for all remaining units $u \in \mathbb{U}$.

To be able to discern between different representations of the same unit, e. g., meters, inches, and miles, these can be encoded as separate units, all of the same dimension “length”. It was just this mismatch between the metric and the imperial system of units that led to the loss of the Mars Climate Orbiter [7]. But a mismatch of the units of a model may not only be detected automatically, but in some cases even be resolved automatically, as some units can be converted into each other, e. g., meters and miles. For this automatic conversion, the function $\tau : \mathbb{U} \times \mathbb{U} \rightarrow component$ is introduced.

The component always has the same signature $i \mapsto o$ with types $i, o : \forall X \sqsubseteq Float. X$. Thus the conversion components can be type-checked with the same rules as the rest of the model. If two units cannot be converted, the function returns \perp .

The function τ can also be interpreted as a *conversion table*¹, as illustrated in an example with different units for the dimension length:

	m	cm	inch	mile
m	\perp	$x * 100$	\perp	$x/1609$
cm	$x/100$	\perp	\perp	\perp
inch	\perp	\perp	\perp	$x/63.360$
mile	$x * 1609$	\perp	$x * 63.360$	\perp

The conversion table τ does not have to be fully specified, since the conversion from cm to inch can be derived from the table above over m and mile. To achieve that, the conversion table can be interpreted as a *conversion graph* by creating a node for each table entry and connecting the nodes whenever the entry of the conversion table is not \perp . Thus the adjacency matrix A for this graph is given as:

$$A(\mu_1, \mu_2) = \begin{cases} 0 & \text{if } \tau(\mu_1, \mu_2) = \perp \\ 1 & \text{otherwise} \end{cases}$$

For the function τ to be consistent, it is required that, (1) if there is a path between any two units μ_1 and μ_2 , there must also be a path between μ_2 and μ_1 , (2) identical units are not converted, i. e., $\forall \mu. \tau(\mu, \mu) = \perp$ and (3) the conversion components do not contain any delays. Note that the

¹To make the table more concise, its components are represented as mathematical functions, rather than SCL components.

adjacency matrix does not have to be symmetric to satisfy these constraints.

Lemma 5 *Two units μ_1, μ_2 are convertible if there exists a directed path in the adjacency matrix A from μ_1 to μ_2 .*

Based on the adjacency matrix, the sets of convertible components could also be defined over the connected components of the conversion graph: Two units μ_1, μ_2 are convertible, if both belong to the same connected component of the graph defined by the adjacency matrix A .

The concept of units is not limited to the conversion based on scalar functions, as in [8], but may also be applied to arbitrary conversion functions, e. g., for Celsius, Kelvin and Fahrenheit:

	C	F	K
C	\perp	$(x - 32)/1,8$	$x - 273$
F	$(x * 1,8) + 32$	\perp	\perp
K	$x + 273$	\perp	\perp

Furthermore, the units are not restricted to physical units but may also be used, e. g., to convert between different currencies. Thus the choice of units depends largely on the application domain.

In contrast to [8] the units are treated separately from the type system. That is because units (in contrast to types) do not have an influence on the system behaviour at runtime. They are only used to check the consistency of the model. Thus we do not want to argue about them in the operational semantics and thus also in the safety proof.

The inference of units in a model is similar to the type inference. Thus the same notation is used as for the type system: In the unit context Δ port p has unit μ is written as $\Delta \vdash p : \mu$. Again for the different entities of an SCL model, different rules are required, the rules (T-port) and (T-perm) apply here as well. To be able to infer the units in a model, there are several operations required on μ . These are based on [8]. Let $n \in \mathbb{Z}$ in:

$$\begin{aligned} \mu_1 = \mu_2 & \iff \forall u \in \mathbb{U}. \mu_1(u) = \mu_2(u) \\ \mu_1 = \mu_2 \cdot \mu_3 & \iff \forall u \in \mathbb{U}. \mu_1(u) = \mu_2(u) + \mu_3(u) \\ \mu_1 = \mu_2^n & \iff \forall u \in \mathbb{U}. \mu_1(u) = n \cdot \mu_2(u) \\ \text{unitless}(\mu) & \iff \forall u \in \mathbb{U}. \mu(u) = 0 \end{aligned}$$

The test for equality $\mu_1 = \mu_2$ is required to check the consistency of the model. The addition of units $\mu_1 \cdot \mu_2$ is required for the mathematical multiplication and division operators, i. e., when two values are multiplied, their units are added. For the division operator units need to be inverted μ^{-1} . For the square root operator every unit will be divided by 2, i. e., μ^2 . The predicate $\text{unitless}(\mu)$ is true if a port does not have any units.

The blocks in SCL are polymorphic in the sense that they can be used for any unit. There are only restrictions on

$$\begin{array}{c}
\frac{\Delta \vdash p_a, p_b, p_c : \mu}{\Delta \vdash plus(p_a, p_b \rightsquigarrow p_c)} \quad (\text{D-plus}) \\
\frac{\Delta \vdash p_a : \mu_a \quad \Delta \vdash p_b : \mu_b \quad \Delta \vdash p_c : \mu_a \cdot \mu_b}{\Delta \vdash mult(p_a, p_b \rightsquigarrow p_c)} \quad (\text{D-mult}) \\
\frac{\Delta \vdash p_a : \mu_a \quad \Delta \vdash p_b : \mu_b \quad \Delta \vdash p_c : \mu_a \cdot \mu_b^{-1}}{\Delta \vdash div(p_a, p_b \rightsquigarrow p_c)} \quad (\text{D-div}) \\
\frac{\Delta \vdash p_a : \mu^2 \quad \Delta \vdash p_b : \mu}{\Delta \vdash sqrt(p_a \rightsquigarrow p_b)} \quad (\text{D-sqrt}) \\
\frac{\Delta \vdash p_a, p_b : \mu}{\Delta \vdash eq(p_a, p_b \rightsquigarrow p_c)} \quad (\text{D-eq}) \\
\frac{\Delta \vdash p_t, p_e, p_r : \mu}{\Delta \vdash ite(p_c, p_t, p_e \rightsquigarrow p_r)} \quad (\text{D-if}) \\
\frac{\Delta \vdash p_a, p_b, v : \mu}{\Delta \vdash delay((p_a \rightsquigarrow p_b), v)} \quad (\text{D-delay}) \\
\frac{\forall c \in C. \Delta \vdash c \quad \forall c \in ch. \Delta \vdash c}{\Delta \vdash component(S, ch, C, \Phi)} \quad (\text{D-net}) \\
\frac{\Delta \vdash p_1 : \mu_1 \quad \Delta \vdash p_2 : \mu_2 \quad \mu_1 \neq \mu_2 \rightarrow convertible(\mu_1, \mu_2)}{\Delta \vdash chan(p_1 \rightsquigarrow p_2)} \quad (\text{D-chan})
\end{array}$$

Figure 6. Unit inference rules of SCL

the allowed combination of units. The rules are depicted in Fig. 6.

The *plus* block only makes sense, if all ports have the same units. The rule for *minus* is identical to (D-plus) modulo renaming. At the *mult* block, the units have to be added, at the *div* block subtracted as defined in rules (D-mult) and (D-div). At the *sqrt* block (D-sqrt), the exponents of the units of the input port must be divisible by two, since only integral numbers are allowed as exponents.

At the *eq* block, it does not make sense to compare values with different units (D-eq).

Analogue to the typing rules, a component is unit-consistent, if its channels and subunits are unit-consistent (D-net). The conversion of units is done at the channels, connecting ports with different units (D-chan). The function $convertible(\mu_1, \mu_2)$ is true, iff there exists a path from μ_1 to μ_2 in τ .

Since units only make sense on numerical types, all ports with units must have a subtype of Float. Thus for all ports with type $\Gamma \vdash p : T$ and unit $\Delta \vdash p : \mu$ the following property must hold:

$$\neg \text{unitless}(\mu) \rightarrow \Gamma \vdash T \sqsubseteq \text{Float}$$

At compile time, after types and units have been checked and inferred successfully, this property is verified by the type checker. This provides the ability to use units not only on a single type, as in [8], but also on any subtype of Float.

The information about units is removed from the model at compile time, since it is no longer needed. The only thing that must remain in the model are components that convert between the different units. These components are generated automatically at compile time. Since the unit conversion only occurs at channels with different units, only these channels need to be modified for unit conversion. These components are generated by replacing the original channel




Figure 7. Conversion from cm to inch

and linking the conversion components along the shortest path between *cm* and *inch* in the conversion graph. The example in Fig. 7 illustrates the components inserted for this the conversion.

5 Examples

To illustrate the static interfaces of SCL, a case study is presented. Here the *remaining range* function was chosen. It is a part of the on-board computer of modern vehicles. Based on the mass of fuel used over a certain distance and the remaining fuel mass, it estimates the remaining range of the vehicle.

As this calculation is not always exact and to have some safety if the driver is misjudging the distance to the next gas station, usually a safety of 5 liter is deducted from the amount of remaining fuel. So that, if the remaining range reaches zero, there will still be about 5 liter of fuel in the tank. A SCL model of the remaining range function is presented in Figure 8. The internal ports in this figure are numbered from p_1 to p_{21} .

5.1 Units

The units of the external ports and the constants have been specified by the user. In this example, these are:




Figure 8. The remaining range case study

Port	μ
<i>remaining_fuel</i>	$[l]$
<i>used_fuel</i>	$[l]$
<i>distance</i>	$[m]$
<i>remaining_range</i>	$[m]$
p_7	$[l]$
p_{11}	$[l]$

$$\begin{aligned}
 \mu_3 &= \mu_1 \cdot \mu_2^{-1} = \\
 &= \mu_{used_fuel} \cdot \mu_{distance}^{-1} = \\
 &= [l] \cdot [m]^{-1} = \\
 &= [l] \cdot [m^{-1}] = \\
 &= [lm^{-1}]
 \end{aligned}$$

The units of p_6 can be derived in the same way. The Inference of $\mu_{14} = [l]$ is not shown here.

To illustrate the inference of units, the ports p_3, p_4, p_5 and p_6 are regarded in detail, since in this example only the division blocks change the units. At the other blocks, the units are simply propagated from the input ports to the output ports.

The units of port p_3 are derived by $\mu_3 = \mu_1 \cdot \mu_2^{-1}$, as defined in rule (D-div). The ports p_1 and p_2 must have the same units as the ports *used_fuel* and respectively *distance* as these are connected by a channel (D-chan). Therefore the units of p_3 are inferred to lm^{-1} :

$$\begin{aligned}
 \mu_6 &= \mu_4 \cdot \mu_5^{-1} = \\
 &= \mu_{14} \cdot \mu_3^{-1} = \\
 &= [l] \cdot [lm^{-1}]^{-1} = \\
 &= [l] \cdot [l^{-1}m] = \\
 &= [m]
 \end{aligned}$$

The channel between ports p_6 and *remaining_range* are consistent considering their units, as both have unit $[m]$.

5.2 Types

In this example, the user has only specified the types of the external ports and the constants. The types of the remaining ports are inferred by the type checker.

$$p_7 : \text{Float}$$

$$p_{11} : \text{Float}$$

In the following, the types of the ports p_8, p_9, p_{10} of the minus block are examined in more detail. The rule (T-minus) states, that all ports must have the same type $p_8, p_9, p_{10} : T$ and that a minus operator must be defined on that type $T \sqsubseteq \{\exists X, \{value : X, minus : X, X \rightarrow X\}\}$

As there exists a channel between ports $p_7 : \text{Float}$ and p_8 and according to rule (T-chan) T must be a supertype of Float, i. e. $\text{Float} \sqsubseteq T$. With rule (U-trans) follows $\text{Float} \sqsubseteq \{\exists X, \{value : X, minus : X, X \rightarrow X\}\}$. This can be shown by applying the the definition of Float and the rules (U-ex) and (U-rcd).

Another example are the ports p_{17} and p_{18} . From the rules (T-gt) and (T-if) follows that $p_{17}, p_{18} : \text{Bool}$. With rule (T-chan) follows that $\text{Bool} \sqsubseteq \text{Bool}$, which is can be concluded from rule (U-refl).

The types of the remaining ports of this example can be inferred in the same way. To instantiale the existential types throughout the model, here Float can be used as it provides all required operators.

6 Conclusions

The framework of static compatibility presented in this paper comprises three different aspects: (1) data type checking, (2) unit checking, and (3) automatic data type and unit inference. This framework enables users to define syntactic interfaces of components by abstracting from concrete types and thereby allows a component to be (re-) used in several concrete contexts (i. e., polymorphic signatures). A concluding abstract example of type and unit inference is depicted in Fig. 9, assuming that the environment provides the appropriate types and units.

This general framework can easily be adopted to other component-based languages exhibiting different syntax and semantics, such as AutoFocus or MATLAB/Simulink, for instance. Notably, our notion of component compatibility can be checked statically, i. e., at compile-time, and therefore does not negatively affect runtime efficiency. Moreover, we have shown that this scheme does not leave space for subtle type-errors that may occur at runtime, and which are often experienced with standard component-based design and modelling tools that lack mechanisms of type inference and polymorphism.

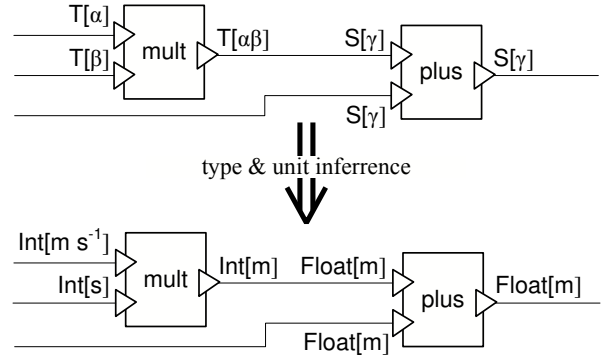


Figure 9. Type and unit inference

In the present form, our approach can be realised with unification in only linear time w. r. t. the size of the model under scrutiny. However, if additional types, such as structured types (tuples, lists, etc.) are introduced to the model, the algorithm will be exponential in the worst-case. This is due to the fact that a component in a model could take as input, say, a tuple, and return, say, a tuple of tuples, and so forth.

Acknowledgements. The authors thank Stefan Berghofer for helpful discussions regarding the subtleties of type systems, and for his comments on draft versions of the paper.

References

- [1] ClawZ—The semantics of Simulink diagrams. Whitepaper, 2003. www.lemma-one.com/clawz_docs.
- [2] A. Benveniste, P. Caspi, S. Edwards, N. Halbwachs, P. Le Guernic, and R. de Simone. The synchronous languages 12 years later. *Proceedings of the IEEE*, 91(1), January 2003.
- [3] G. Berry. *Proof, Language and Interaction: Essays in Honour of Robin Milner*, chapter The Foundations of Esterel. MIT Press, 2000.
- [4] M. Broy, F. Huber, and B. Schätz. AutoFocus – Ein Werkzeugprototyp zur Entwicklung eingebetteter Systeme. *Inform., Forsch. Entwickl.*, 14(3):121–134, 1999.
- [5] P. Caspi, A. Curic, A. Maignan, C. Sofronis, S. Tripakis, and P. Niebert. From simulink to SCADE/Lustre to TTA: a layered approach for distributed embedded applications. In *ACM SIGPLAN Conf. on Languages, compilers, and tools for embedded systems*. ACM Press, 2003.
- [6] L. de Alfaro and T. A. Henzinger. Interface automata. In *Proc. 8th European software engineering conference*, New York, NY, USA, 2001. ACM Press.
- [7] E. E. Euler, S. D. Jolly, and H. H. Curtis. The failures of the mars climate orbiter and mars polar lander: A perspective from the people involved. In *Proc. Guidance and Control*. American Astronautical Society, 2001.
- [8] A. J. Kennedy. Relational parametricity and units of measure. In *Proc. ACM SIGPLAN-SIGACT symposium on Principles of prog. lang.* ACM Press, 1997.

- [9] K. Knight. Unification: a multidisciplinary survey. *ACM Comput. Surv.*, 21(1):93–124, 1989.
- [10] C. Kühnel, A. Bauer, and M. Tautschnig. Compatibility and reuse in component-based systems via type and unit inference. In *Proceedings of the 33rd EUROMICRO CONFERENCE on Software Engineering and Advanced Applications*, 2007.
- [11] E. A. Lee and Y. Xiong. System-level types for component-based design. In *EMSOFT'01*. Springer, 2001.
- [12] NIST. The NIST reference on Constants, Units and Uncertainty. physics.nist.gov/cuu/Units.
- [13] B. C. Pierce. *Types and programming languages*. MIT Press, Cambridge, MA, USA, 2002.
- [14] G. D. Plotkin. The origins of structural operational semantics. *Journal of Logic and Algebraic Programming*, 60-61:3–15, 2004.