

# Scaleness/nesT: Type Specialized Staged Programming for Sensor Networks

Peter Chapin  
University of Vermont  
pchapin@cs.uvm.edu

Christian Skalka\*  
University of Vermont  
skalka@cs.uvm.edu

Scott Smith  
The Johns Hopkins University  
scott@cs.jhu.edu

Michael Watson  
University of Vermont  
mpwatson@cs.uvm.edu

## Abstract

Programming wireless embedded networks is challenging due to severe limitations on processing speed, memory, and bandwidth. Staged programming can help bridge the gap between high level code refinement techniques and efficient device level programs by allowing a first stage program to specialize device level code. Here we introduce a two stage programming system for wireless sensor networks. The first stage program is written in our extended dialect of Scala, called Scaleness, where components written in our type safe dialect of nesC, called nesT, are composed and specialized. Scaleness programs can dynamically construct TinyOS-compliant nesT device images that can be deployed to motes. A key result, called cross-stage type safety, shows that successful static type checking of a Scaleness program means no type errors will arise either during programmatic composition and specialization of WSN code, or later on the WSN itself. Scaleness has been implemented through direct modification of the Scala compiler. Implementation of a staged public-key cryptography calculation shows the sensor memory footprint can be significantly reduced by staging.

**Categories and Subject Descriptors** D.3.4 [Programming Languages]: Processors—Compilers

**Keywords** Staged programming; Scala; nesC

## 1. Introduction

A wireless sensor network (WSN) is a network of small nodes, also called *motes*, equipped with sensors or actuators, and that communicate with each other via short range radio links. Programming WSNs is challenging because the nodes are severely resource constrained in terms of memory and processor speed. This paper describes a programming language designed to support the automatic generation of runtime-efficient code for WSN nodes. The language enables *dynamic specialization* of node code on a nearby hub or other more resource-rich device, allowing adaptation to properties of a node’s deployment environment such as neighborhood characteristics, network interference factors, etc.

\*This author’s work was supported by a YIP grant from the Air Force Office of Scientific Research (AFOSR).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

GPCE '13, October 27–28, 2013, Indianapolis, Indiana, USA.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-2373-4/13/10...\$15.00.

<http://dx.doi.org/10.1145/2517208.2517217>

Our system supports dynamic generation of programs for the TinyOS operating system [16], a popular platform for WSNs. It features programming abstractions for specializing WSN code, allowing on-the-fly adaptation to current WSN deployment conditions. We use a restricted form of *staging* [8, 32, 33] to achieve well founded dynamic program generation. *First stage* code is written in an extended version of Scala [30], called Scaleness. Scaleness program execution yields a residual *second stage* WSN node program written in nesT, a variant of the popular nesC programming language [12] with a stronger type checking analysis. The second stage program is constructed from module components treated as first class values, which may be *type and value specialized* during the course of first stage computation to yield more compact and efficient code. A code rewriting strategy in the implementation transforms nesT code into nesC code, which can be compiled using standard TinyOS tools.

While staging is well-studied and has been explored in a WSN context [24], our work is novel in that we achieve stronger static safety guarantees than previous work. At the point of Scaleness program compilation, our compiler can statically verify that any nesT program produced by the Scaleness runtime will be statically type-safe when deployed and run on a network node, even if module parameters are specialized during the course of nesT module composition. We call this property *cross-stage type safety*, which has been previously studied in a foundational language context [20]. In this paper we apply these concepts to the more practical Scaleness/nesT language, and illustrate how they support the implementation of efficient real-life WSN applications.

### 1.1 Application Setting and Contributions

The diagram in Fig. 1 provides an overview of the Scaleness/nesT language architecture. Scaleness source code is compiled in a modified Scala compiler to Java bytecode, and run in a standard JVM. At runtime this Scaleness program may generate nesT code, which is subsequently rewritten to nesC and compiled using the standard TinyOS compiler. The resulting image can then be installed on nodes in a WSN.

Another interesting feature of our intended application setting, captured in Fig. 1, is the physical platform on which different elements of the Scaleness/nesT “workflow” may be executed. Scaleness source code will typically be compiled in the lab, prior to deployment. There are two distinct deployment scenarios where compiled bytecode execution, TinyOS image generation, and mote (re)programming (the rightmost two boxes in Fig. 1) can occur. Clearly these activities can take place in the lab, where WSN motes can be easily imaged over e.g. USB connections prior to deployment. But the more interesting scenario we aim to support is generation of TinyOS images on a “hub” device *in situ*, and then to automatically reprogram WSN nodes over the air (OTA) from that hub.

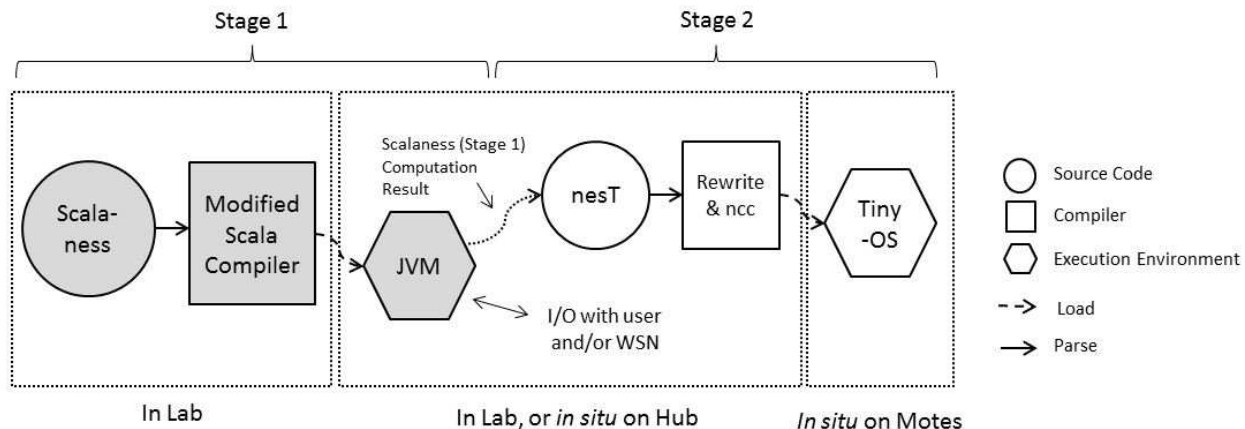


Figure 1. Scalanness/nest Compilation and Execution Model

In WSN applications such as our Snowcloud snow telemetry system [11, 26], sensor motes report data to higher powered hubs, pictured in Fig. 2. The hub device in the figure uses a low-powered embedded Linux platform that supports a JVM, and is in direct radio communication with the sensor network via a physically attached mote. These types of devices are readily available and cheap. Such a system can execute compiled Scalanness code, and compile and deploy generated TinyOS images to nearby motes. Since the hub is in communication with the network, Scalanness/nest is uniquely positioned to evolve network behavior based on a global view of observed data, a technique called *backcasting* [34] when used specifically for network control. In this context the benefit of cross-stage type safety is clear: type-checked Scalanness compilation in the lab ensures type safety of bytecode execution on the hub, and type-safety of dynamically generated TinyOS image execution on the WSN. Manual correction of type errors in generated TinyOS images in this scenario is infeasible since these systems run automatically in remote settings.




Figure 2. A Sensor Node (L,C) and Hub Device (R).

**Paper Outline.** The main contributions of the work presented here are the specification, implementation, and prototype application of the Scalanness and nest languages, including their syntax, semantics, and type analysis. In Sect. 2 we summarize Scalanness/nest via discussion of an extended example. Formal specifications of the nest and Scalanness languages are presented in Sect. 3 and Sect. 4, respectively. Their semantics and type theory possess several novel and challenging features, which we show in Sect. 5 are grounded in principles studied in a previous foundational setting [20]. We describe our implementation and present

an extended example application of our system to resource access control in WSNs in Sect. 6 and Sect. 7, along with some empirical results demonstrating efficiency benefits of our approach. We conclude with remarks on related work in Sect. 8

## 2. An Example: Authenticated Messaging

In this section we provide a high-level overview of Scalanness/nest usage and applications via an example. (This example is written in DScalanness/Dnest, a simplified formalization of the implemented Scalanness/nest, which is defined later in this paper.) The example illustrates both the type and value specializations that can occur in our system. This example needs to be very small to fit in this overview which means it alone is not complete evidence of the advantage of mote code staging, but it should point the reader to the promise of the approach.

**Program description.** To illustrate type specialization, we refine address size. It is well-known that minimizing address space size in WSN message packets can obtain significant energy savings by reducing message sizes, since each bit of transmission is known to consume energy similar to 800 instructions [22]. However, WSNs are “ad hoc” precisely in the sense that positions and densities of nodes in space are unpredictable, so “minimal” address space is an environmental property, where minimality may need to be determined *in situ*.

To illustrate value specialization, we define a Dnest code template that can be instantiated with specific session keys for secure communication in a WSN. We imagine that the template is instantiated on high powered hub or lab device, where session keys are generated. In previous work it has been shown how symmetric key signatures can be used to support language based resource authorization in WSNs [4, 5]. In particular, communication between security domains in a WSN is mediated by credentials implemented as keys, and nodes lying at domain frontiers can use different keys to send (to the other domain) and receive (from the other domain) over secured link layer channels. Since it is unpredictable where nodes will be physically distributed in space, appropriate keys for each node need to be established *in situ*. Defining node functionality using generic code that must be instantiated with specific keys allows adaptation to a deployment environment, and allows expensive computations for establishing session keys to be offloaded from the WSN to a higher powered device. Experience with an actual implementation of this application is discussed Sect. 7.

**The Code.** To distinguish Scalness and nesT code in examples we will use a darker font for Scalness code and a lighter font for nesT code, and line numbers for reference. We begin with the definition of a parameterized type `msgT(t)` using the Scalness `abbrvt` binder, where an instance `msgT( $\tau$ )` denotes the ground type obtained by substituting  $\tau$  for `t` in the definition of `msgT`.

```
1 abbrvt msgT(t) = { src : t; dest : t; data : uint8[] };
```

Next, we define a type `radioT`, which is the type of nesT modules that provide an API to the radio.

```
2 abbrvt radioT = < mt  $\preceq$  msgT(uint) >
3   { export radio_x(mt*);
4     import handle_radio_r(mt*); };
```

The nesT module language is a simplified version of the nesC component language. In this example, any module of type `radioT` exports a `radio_x` function for sending messages, and imports a `handle_radio_r` function that allows received messages to be handled in a user-defined manner. Both functions take message references as arguments<sup>1</sup>. Furthermore, the module is parameterized by the type of messages `mt`, where the address type is upper-bounded by 32-bit unsigned integer. Thus, any module of type `radioT` can be dynamically specialized to a 32, 16, or 8 bit address space by type instantiation. Module type parameters are always defined with brackets `< ... >`.

Now we define another type `commT` which is the type of modules providing a QOS layer over a specialized radio.

```
5 abbrvt commT = (mt  $\preceq$  msgT(uint)) o
6   < >
7   { export send(mt*);
8     import handle_receive(mt*); };
```

Although this type is also parameterized by a bounded message type `mt`, as is `radioT`, the parameterization is subtly different syntactically and semantically, since `commT` expects a program context where the radio has been specialized. Thus, in `commT`, `mt` is understood as being “some” type with an upper bound of `msgT(uint)` which occurs in the module signature, whereas the module itself has no parameters to be instantiated—note the empty instance parameter brackets `<>` in the module type after the `o` delimiter. This sort of type is needed in the presence of *dynamic type construction*, a useful Scalness feature we exemplify below.

Next we define modules for sending and receiving messages that provide a layer of authentication security over the radio.

```
9 authSend = < mt  $\preceq$  msgT(uint); sendk : uint8[], >
10   { import radio_x(mt*);
11     export send(m : mt*)
12       { radio_x(AES_sign(m, sendk)); } };
13
14 authRecv = < mt  $\preceq$  msgT(uint); recvk : uint8[] >
15   { import handle_recv(mt*);
16     export handle_radio_r(m : mt*)
17       { if AES_signed(m, recvk)
18         handle_recv(m); } };
```

Observe that in the implementation of `send` in module `authSend`, messages are signed with a key `sendk`, whereas when messages are received they must be signed with a possibly different key `recvk` before being passed on to the user’s receive handler, as specified in module `authRecv`. These modules are parameterized by a message type `mt`, and also the `sendk` and `recvk` key values.

To generalize a technique for composing these modules with a radio to yield a module of type `commT`, that is abstract wrt neighborhood sizes, radio implementations, and session key material, we define the Scalness `authSpecialize` function as follows:

```
19 def authSpecialize
20   (nmax : uint16, radioM : radioT, keys : uint8[][] ) : commT {
```

```
21   typedef adt  $\preceq$  uint = if (nmax  $\leq$  256) uint8 else uint16;
22   val sendM = authSend(msgT(adt); keys[0]);
23   val recvM = authRecv(msgT(adt); keys[1]);
24   (sendM  $\times$  radioM(msgT(adt)))  $\times$  recvM;
25 }
```

The first-class status of nesT modules in Scalness is apparent here. On line 20 the function is specified to take a module parameter `radioM` of type `radioT` among its arguments, and to return a module of type `commT` as a result. It also takes an array of keys as an argument, and on lines 22 and 23 it instantiates `sendMsg` and `recvMsg` with the keys in the array. It also uses the type `adt` in the instantiations, which in line 21 is dynamically constructed on the basis of the input variable `nmax` which defines the needed address space size and bound using the Scalness `typedef` construct. This illustrates a key novelty of our system, the ability to *dynamically* set a type to use on a mote based on a decision made in the Scalness runtime. Since the value of `nmax` cannot be statically determined, the type analysis only knows that `adt` is some subtype of `uint`. Finally, on line 24 the instantiated radio module is composed with the instantiated send and receive modules via the Scalness  `$\times$`  operator. The semantics of module composition here is standard [2]; in a composition aka wiring  $\mu_1 \times \mu_2$ , the exports of  $\mu_2$  are connected to imports of  $\mu_1$ . The function result is a module of type `commT`.

To obtain a module defining a mote OS image in a program context where neighborhood size is known, a radio implementation has been provided, and session keys have been computed. We can then compose the results of an `authSpecialize` function with modules specifying top-level message send and receive behaviors, and a main application entry point as follows (here we assume it is known that address sizes can be limited to 8 bits, so `nmax < 256`). At line 30 a closed module is defined and a binary mote image can be produced by a call to `image`.

```
26 appMR =
27   < > { export handle_recv(m : msgT(uint8*)) {...} };
28 appM =
29   < > { import send(msgT(uint8*)) ; export main() {...} };
30 image(appM  $\times$  (authSpecialize(nmax, radioM, keys)  $\times$  appMR));
```

In DScalness, `image` is an assertion that its argument is a *runnable* module, with no unresolved parameters or imports. In the Scalness implementation, this is the point where nesT source code is actually generated. Successful Scalness/nesT type checking (which occurs during stage 1 compilation as per Fig. 1) statically guarantees that specialized code generated at the point of `image` will run in a type-safe manner when it is eventually loaded and run on a mote.

### 3. The nesT Language Distilled

In this section we summarize a *Distilled* version of nesT, called `DnesT`, that isolates novel elements of nesT, specifically parametric types, subtyping, type safety, and modules. `DnesT` serves as a formal specification for the nesT implementation—given the novel type theory a specification is crucial as a guide for the implementation. Due to limited space we summarize only the top-level structure of `DnesT` modules and our type checking algorithm, in order to focus more on the more central technical issues of module composition, instantiation, and typing at the Scalness level.

The goal of nesT is to be a type-safe variant of nesC, and `DnesT` serves as the specification for how type safety is achieved. Our approach is another species of “safe C” language design projects such as [29]. In particular, in `DnesT` all array bound accesses are checked at run-time, and pointer arithmetic and casting are restricted to safe forms only. We have developed a new type checking algorithm that incorporates subtyping, which supports bounded type parameters in `DnesT` module definitions and a more accurate static analysis of Scalness code in the presence of type construction and nesT module instantiation.

<sup>1</sup>For brevity the return type on all commands is omitted. In all cases it is the TinyOS error type `error_t`

$\varsigma, \tau$	$::= t \mid \top \mid \mathbf{uint8} \mid \mathbf{uint16} \mid \mathbf{uint} \mid \mathbf{uninit} \mid \{\bar{l} : \bar{\tau}\} \mid \tau[] \mid \tau^*$	types
$T$	$::= \bar{l} \preceq \bar{\tau}$	type parameters
$V$	$::= \bar{x} : \bar{\tau}$	value parameters
$c$	$::= f(V) : \tau = \{e\}$	command definition
$s$	$::= f(V) : \tau$	command signature
$\iota$	$::= \bar{s}$	imports
$\xi$	$::= \bar{c}$	exports
$\varepsilon$	$::= \bar{s}$	export types
$d$	$::= \tau x = e \mid \tau x = [] \bar{e}[] \mid \tau x = \{\bar{l} = \bar{e}\} \mid c$	declarations
$\mu$	$::= \langle T; V \rangle \{ \iota; \bar{d}; \xi \}$	module definitions
$\mu\sigma$	$::= \langle T; V \rangle \{ \iota; \varepsilon \}$	module signatures

Figure 3. Syntax of DnesT Types and Modules

### 3.1 Syntax and Semantics of DnesT

Module definitions rely on a notion of lists aka sequences of syntactic entities, so we begin with a definition of relevant notation.

**Notation and identifiers.** Sequences are notated  $x_1, \dots, x_n$ , and are abbreviated  $\bar{x}$ ;  $\bar{x}_{(i)}$  is the  $i$ -th element,  $\emptyset$  denotes the empty sequence, and  $|\bar{x}|$  is the size. We write  $x \in \bar{x}$  to denote membership in sequences, and  $x\bar{x}$  denotes a sequence with head  $x$  and tail  $\bar{x}$ . We denote append as  $\bar{x}\bar{y}$ . For relational symbols  $R \in \{\preceq, =, \cdot\}$ , we use the abbreviation:  $\bar{x}R\bar{y} = x_1 R y_1, \dots, x_n R y_n$ . So for example,  $\bar{x} : \bar{\tau} = x_1 : \tau_1, \dots, x_n : \tau_n$ . We will use metavariable  $f$  (of set  $\mathcal{F}$ ) for function names,  $l$  (of set  $\mathcal{L}$ ) for field names,  $x$  (of set  $\mathcal{V}$ ) for term variables,  $t$  (of set  $\mathcal{T}$ ) for type variables.

**Module syntax.** The syntax of DnesT modules is defined in Fig. 3. Modules  $\mu$  are written  $\langle T; V \rangle \{ \iota; \bar{d}; \xi \}$  with  $T$  and  $V$  being generic type and term parameters,  $\bar{d}$  being module scope identifier declarations, including function definitions, and  $\iota$  and  $\xi$  being imports and exports. In Sect. 2 and elsewhere we use the keywords `import` and `export` as sugar indicating this categorization.

All type parameters are assigned an upper bound, and term parameters are explicitly typed. Imports and export types are sequences of imported and exported command type signatures. Exports are sequences of command definitions. Exports are defined in terms of expressions  $e$ , the syntax of which we omit here for brevity. Declarations  $\bar{d}$ ; are a sequence of typed variable declarations. Base values, arrays (in brackets  $[\cdot]$ ), structs (in braces  $\{\cdot\}$ ), and commands may all be declared, and the scope of declared variable names is restricted to the module. Declarations are important to include in DnesT, as they support serialization of value parameters during Scalness instantiation as we describe in Sect. 4.3.

While we have elided the specifics of DnesT syntax from this shortened presentation, we now give a high-level summary of its largely standard features. Expressions include standard C-like conditional branching, looping, sequencing of expressions, function calls, arrays, structs, numeric base data types and basic arithmetic operations. As in nesC, no dynamic memory allocation is possible; all memory layout is established by static variable declarations. DnesT disallows pointer arithmetic, to support stronger type safety guarantees. Type casting and array access have run time checks imposed: types may never be cast to a pointer, and array accesses are always checked to be in bounds at runtime. As in nesC, DnesT includes a `post` operation for posting tasks, although we make no syntactic distinction between tasks and commands. The meaning of `post` corresponds to the “run-to-completion” model of TinyOS tasks. Interrupts are omitted from DnesT since they do not significantly affect the typing issues we are concerned with here.

**Module semantics** A “runnable” module – one without imports or generic parameters – is the DnesT model of a node OS image. The declarations in the module defines a *load sequence* establishing an initial machine configuration, and the application entry point is defined in a required command `main`.

DEFINITION 3.1. A module of the form  $\langle \emptyset; \emptyset \rangle \{ \bar{d}; \xi \}$ , where `main()` : `uninit`  $\in \xi$ , is called runnable.

This model is consistent with nesC, where an application is defined as a top-level component that establishes an initial configuration through variable declarations, and requires user definition of an entry point (an event handler called `Booted`). Formally speaking, type safety in nesT is a dynamic property of runnable modules.

### 3.2 Type Checking and Subtyping

The type system for DnesT combines a standard procedural language typing approach with subtyping techniques adapted from previous foundational work [13, 20].

At the heart of our system is a decidable subtyping judgment  $T \vdash \tau_1 \preceq \tau_2$ , where  $T$  is a *coercion* and defines a system of upper bounds for type variables. This establishes a subtype ordering on base types, and also allows for width subtyping of records. The relation is defined in Fig. 4. Algorithms for deciding the relation and integrating it with dynamic type construction and other Scalness (stage 1) type features was a central topic of [20].

REFLS $T \vdash \tau \preceq \tau$	TOPS $T \vdash \tau \preceq \top$	TRANS $\frac{T \vdash \tau_1 \preceq \tau_2 \quad T \vdash \tau_2 \preceq \tau_3}{T \vdash \tau_1 \preceq \tau_3}$
UINTS $T \vdash \mathbf{uint8} \preceq \mathbf{uint16} \preceq \mathbf{uint}$	STRUCTS $\frac{T \vdash \tau_1 \preceq \tau_3}{T \vdash \{l_1 : \tau_1 \uplus l_2 : \tau_2\} \preceq \{l_1 : \tau_3\}}$	

Figure 4. Subtyping Rules

The type checking algorithm for DnesT expressions is a combination of standard procedural type systems and standard subtyping systems. Module typing is obtained by type checking module exports, using a coercion obtained from the module type parameters and a typing environment obtained from a combination of module value parameters, imports, and variable type declarations. A valid module type checking judgment is written as:

$$\langle T, V \rangle \{ \iota; \bar{d}; \xi \} : \langle T, V \rangle \{ \iota; \varepsilon \}$$

Where  $\varepsilon$  is just the type signatures of  $\xi$ , and each of the command bodies in  $\xi$  is proven to respect its type signature.

EXAMPLE 3.1. The module `authSend` defined in Sect. 2 code line 9 can be assigned the following type in DnesT:

```
< mt ≲ msgT(uint); sendk : uint8[] >
{ import radio_x(mt*), export send(mt*) }
```

## 4. The Scalness Language Distilled

Scalness serves as the language for nesT module composition in the same manner as nesC configurations serve to compose nesC modules, but Scalness is a much more powerful metalanguage since modules are treated as a new category of first class values in Scalness. Instantiation, composition (aka wiring), and imaging of modules are defined as operations on module values. Because instantiation of modules with both types and values is allowed, types and values may migrate from the Scalness level to the nesT level after programmatic refinement, realizing a disciplined form of code specialization.

$L$	::=	$\text{class } C(\bar{X} <: \bar{N}) \text{ extends } N \{ \bar{T} \bar{f}; K \bar{M} \}$	<i>classes</i>
$K$	::=	$C(\bar{T} \bar{f}) \{ \text{super}(\bar{f}); \text{this}.\bar{f} = \bar{f}; \}$	<i>constructors</i>
$M$	::=	$T m(\bar{T} \bar{x}) \{ \text{return } e; \}$	<i>methods</i>
$e$	::=	$x \mid e.f \mid e.m(\bar{e}) \mid \text{new } C(\bar{T})(\bar{e}) \mid (N)e \mid$ $1 \mid e.f = e \mid \text{def } x : T = e \text{ in } e \mid$ $\text{abbrvt } X(\bar{X}) = T \text{ in } e \mid$ $\mu \mid e \times e \mid e(\bar{e}; \bar{e}) \mid \text{image } e$	<i>expressions</i>
$T$	::=	$X \mid N \mid T \circ \mu r$	<i>types</i>
$N$	::=	$C(\bar{T})$	<i>class types</i>
$1$	::=	$(p, N)$	<i>references</i>

**Figure 5.** The Syntax of DScalanness

Our goal in this Section is to describe the Scalanness syntax and semantics realized in our implementation. Since Scala is too large to easily formalize, we define here a *Distilled Scalanness*, DScalanness, that extends a core typed object-oriented language to include syntax and semantics for defining and composing DnesT modules. The particular object-oriented core calculus we use is a combination of two Featherweight Java variants: Featherweight Generic Java (FGJ) [17] and Assignment Featherweight Java (AFJ) [27].

#### 4.1 Syntax of DScalanness

The DScalanness language syntax is presented in Fig. 5. We refer the reader to [17, 27] for details on the FGJ and AFJ object oriented calculi, which are represented in the languages of class definitions, constructors, methods, and the first line of expression forms defined in Fig. 5. DScalanness extends these features with a typed variable declaration form  $\text{def } x : T = e_1 \text{ in } e_2$  where the scope of  $x$  is  $e_2$ , a dynamic type construction form  $\text{typedef } x <: T = e_1 \text{ in } e_2$  with similar scoping rules (although this is defined as syntactic sugar in Definition 4.1). For programmer convenience a simple parameterized type abbreviation binder  $\text{abbrvt}$  is also provided. We include DnesT modules  $\mu$  in the DScalanness expression and value spaces: instantiation is obtained via the form  $e_1(\bar{e}_1; \bar{e}_2)$ , where  $\bar{e}_1$  are type parameters and  $\bar{e}_2$  are value parameters. Wiring of modules is denoted  $e_1 \times e_2$ . Imaging of modules, denoted  $\text{image } e$ , ensures that  $e$  computes to a runnable DnesT module.

#### 4.2 Semantics of DScalanness

The semantics of DScalanness is an extension of the semantics of AFJ and FGJ to incorporate DnesT modules and operations. Computations assume a fixed class table  $CT$  allowing access to class definitions via class names, which always decorate an object's type. A *store*  $ST$  is a function from memory locations  $p$  to object representations. Objects are represented in memory by lists of object references  $\bar{1}$ , which refer to the locations of the objects stored in mutable field values. A reference  $1$  is a pair  $(p, N)$  where  $p$  is the memory location of an object representation and  $N$  is the nominal type of the object, including its class name. Hence, given an object reference  $(p, C(\bar{T}))$ , we can access and mutate its fields  $\bar{1} = ST(p)$ , and access and use its methods via the definition  $CT(C)$ .

Following AFJ, the semantics of DScalanness is defined as a *labeled transition system*, where transitions are of the form  $e - \{s = ST, s' = ST'\} \rightarrow e'$ . Intuitively, this denotes that given an initial store  $ST$  and expression  $e$ , one step of evaluation results in a modified store  $ST'$  and contractum  $e'$ . We write  $e \rightarrow e'$  as an abbreviation when the store is not altered.

The primary novelty of DScalanness is the formal semantics of type and module construction. We begin with type construction, which is provided to allow programmers to dynamically con-

struct module type instances. The appropriate behavior is obtained by treating dynamically constructed types as extensions of a basic class of objects, and declarations of DnesT level types via a *typedef* construct as syntactic sugar for ordinary object construction. We define a `LiftableType` class as the supertype of all types of objects that can be used to instantiate a module, and dynamically constructed types are defined as instances of a generic `MetaType` class.

**DEFINITION 4.1.** Any DScalanness class table  $CT$  comprises the following definitions:

```

CT(LiftableType) =
  class LiftableType() extends Object {...}
CT(MetaType) =
  class MetaType(X <: LiftableType) extends Object {...}

```

And we take as given the following syntactic sugar:

```

typedef x <: T = e_1 in e_2  $\triangleq$  def x : MetaType(T) = e_1 in e_2

```

Class type `MetaType` is generalized on a single type variable. For brevity of notation, we define:

$$\text{MetaType}(\bar{T}) \triangleq \overline{\text{MetaType}(\bar{T})}$$

A crucial fact of DScalanness type construction is that any dynamically constructed type cannot be treated as a type at the DScalanness level. This is a more restrictive mechanism than envisioned in our foundational model [20], however it allows us to define DScalanness as a straightforward extension to Scala, especially in terms of type checking.

Module instantiation, shown in Fig. 6, is the only point where specialization of DnesT modules is allowed. Since DScalanness and DnesT are two different language spaces, some sort of transformation must occur when values migrate from DScalanness to DnesT via module instantiation. This *lifting* transformation involves both data mapping and serialization since the process spaces also differ. We aim to be flexible and allow the user to specify how values are lifted and how types are transformed. We only require that lifting and type transformation are coherent, in the sense that the lifting of an object should be typeable at the object's type transformation. We formalize this in the following definition.

**DEFINITION 4.2.** We assume given a relation  $\xrightarrow{\text{lift}}$  which transforms a DScalanness reference  $1$  into DnesT declarations  $\bar{d}$  and expression  $e$ . We also assume given a DScalanness-to-DnesT transformation of types  $[\cdot]$ . To preserve type safety, we require in all cases that  $(p, N) \xrightarrow{\text{lift}} \bar{d}, e$  implies both of the following for some type environment  $G$ :

$$\emptyset, \emptyset \vdash \bar{d} : G \quad \text{and} \quad G, \emptyset \vdash e : \llbracket N \rrbracket$$

The full definition of serialization and an example are given and discussed below in Sect. 4.3.

Module wiring is given a standard component composition semantics. We only allow wiring of instantiated modules, which is consistent with nesC and simpler to implement. In a wiring  $e_1 \times e_2$ , the exports of  $e_2$  are wired to the imports of  $e_1$ . This is specified in the `MODWIRE` rule in Fig. 6, which relies on the following auxiliary definition of operations for combining mappings.

**DEFINITION 4.3 (Special Mapping Operations).** Let  $m$  range over vectors with mapping interpretations, in particular  $T, V, \iota$ , and  $\xi$ . Binary operator  $m_1 \curlywedge m_2$  represents (non-exclusive) map merge, i.e.  $m_1 \curlywedge m_2 = m_1 @ m_2$  with the requirement that  $id \in \text{Dom}(m_1) \cap \text{Dom}(m_2)$  implies  $m_1(id) = m_2(id)$ . The mapping  $m/S$  is the same as  $m$  except undefined on domain elements in set  $S$ , and the mapping  $m \upharpoonright_S$  is the same as  $m$  except undefined on elements not in  $S$ .

$$\text{MODINST} \quad \frac{\mu : \langle \bar{t} \preceq \bar{\tau}; \bar{x} : \bar{\zeta} \rangle \{ \bar{u}; \bar{d}; \xi \} \quad \text{serialize}(\bar{x}, \bar{\zeta}, \bar{1}) = \bar{d}'}{\mu(\langle \bar{p}, \text{MetaType}(\bar{T}); \bar{1} \rangle \rightarrow \langle \rangle \{ \bar{u}; \bar{d}' @ \bar{d}; \xi \} \llbracket \bar{T} \rrbracket / \bar{t})}$$

$$\text{MODWIRE} \quad \frac{\iota = (\iota_1 / \text{Dom}(\xi_2)) @ \iota_2 \quad \bar{d} = \bar{d}_2 @ \xi_2 |_{\text{Dom}(\iota_1)}}{\langle \rangle \{ \bar{u}_1; \bar{d}_1; \xi_1 \} \times \langle \rangle \{ \bar{u}_2; \bar{d}_2; \xi_2 \} \rightarrow \langle \rangle \{ \bar{u}; \bar{d} \vee \bar{d}_1; \xi_1 \}}$$

$$\text{MODIMAGE} \quad \frac{\text{main defined in } \xi}{\text{image}(\langle \rangle \{ \bar{d}; \xi \}) \rightarrow \langle \rangle \{ \bar{d}; \xi \}}$$

**Figure 6.** DScaleness Module Semantics

Finally, the MODIMAGE rule in Fig. 6 shows that imaging it is an assertion requiring its arguments to be a runnable module.

EXAMPLE 4.1. *Given code definitions in Sect. 2 and an invocation `authSpecialize(50, radioM, [[k1, k2]])`, where `radioM : radioT`, and `k1, k2` are keys, the evaluation of the expression on line 24, `sendM × radioM(msgT(adt))`, will evaluate to the following module:*

```
< > ( { import handle_radio_r(msgT(uint8)*); ...;
      export send(m : msgT(uint8)*)
        { radio_x(AES_sign(m, k1)); } }
```

where the elided declarations include a definition of a command `radio_x` imported from `radioM` also with argument type `msgT(uint8)*`.

### 4.3 Serialization and Lifting

Serialization generates a flattened DnesT source code version of a DScaleness object in memory. At the top level, serialization binds the value parameters of a module to the results of flattening, aka lifting, via a sequence of declarations. Here is the precise definition.

DEFINITION 4.4 (Serialization). *Assume given a store  $ST$  which is implicit in the following definitions. We define serialization of DScaleness references as follows, along with an extension of the user defined lifting relation to sequences of references:*

$$\frac{\emptyset \xrightarrow{\text{lift}} \emptyset, \emptyset}{\frac{\bar{1} \xrightarrow{\text{lift}} \bar{d}, \bar{e}}{\text{serialize}(\bar{x}, \bar{\tau}, \bar{1}) = \bar{d} @ \bar{\tau} \bar{x} = \bar{e}} \quad \frac{\bar{1} \xrightarrow{\text{lift}} \bar{d}, e \quad \bar{1} \xrightarrow{\text{lift}} \bar{d}', \bar{e}}{\bar{1} \xrightarrow{\text{lift}} \bar{d} @ \bar{d}', e \bar{e}}}}$$

Although lifting is user defined, a standard strategy is to introduce a new declared variable for each memory reference in the lifted object, and bind the variable to the lifted referent. Hence, lifting will typically be defined recursively. In our implementation, we have adapted a “default” lifting which follows this strategy, and also transforms objects by just transforming the fields into a representative struct, and ignoring methods. We will illustrate this with an example in Sect. 6. We can formally capture the essence of this transformation with the following definitions. It is easy to see that these definitions will satisfy the requirements of Definition 4.2.

EXAMPLE 4.2. *In this example we allow lifting of any object references, and transform the object  $o$  into a structure containing the transformed fields of  $o$ . Methods are disregarded by the transformation. Here is the specification of the type transformation:*

$$\frac{CT(C) = \text{class } C(\bar{X} <: \bar{S}) \text{ extends } N \{ \bar{R} \bar{F}; K \bar{M} \}}{\llbracket C(\bar{T}) \rrbracket = \{ \bar{F} : \llbracket \bar{R}(\bar{T}/\bar{X}) \rrbracket \}}$$

$$\text{MODT} \quad \frac{\mu : \mu\tau \text{ in DnesT type checking}}{\Gamma \vdash \mu : \emptyset \circ \mu\tau}$$

$$\text{MODIMAGET} \quad \frac{\Gamma \vdash e : T \circ \langle \rangle \{ \bar{u}; \varepsilon \} \quad \text{main}() : \tau \in \varepsilon}{\Gamma \vdash \text{image } e : T \circ \langle \rangle \{ \bar{u}; \varepsilon \}}$$

$$\text{MODINSTT} \quad \frac{\Gamma \vdash e : \emptyset \circ \langle \bar{t} \preceq \bar{\tau}_1; \bar{x} : \bar{\tau}_2 \rangle \{ \bar{u}; \varepsilon \} \quad \Gamma \vdash \bar{e}_1 : \text{MetaType}(\bar{T}_1) \quad \Gamma \vdash \bar{e}_2 : \bar{T}_2 \quad \vdash \llbracket \bar{T}_1 \rrbracket \preceq \bar{\tau}_1 \quad \vdash \llbracket \bar{T}_2 \rrbracket \preceq \bar{\tau}_2}{\Gamma \vdash e(\bar{e}_1; \bar{e}_2) : \bar{t} \preceq \llbracket \bar{T}_1 \rrbracket \circ \langle \rangle \{ \bar{u}; \varepsilon \}}$$

$$\text{MODWIRET} \quad \frac{\Gamma \vdash e_1 : T_1 \circ \langle \rangle \{ \bar{u}_1; \varepsilon_1 \} \quad \Gamma \vdash e_2 : T_2 \circ \langle \rangle \{ \bar{u}_2; \varepsilon_2 \} \quad \iota = (\iota_1 / \text{Dom}(\varepsilon_2)) @ \iota_2}{\Gamma \vdash e_1 \times e_2 : T_1 \vee T_2 \circ \langle \rangle \{ \bar{u}; \varepsilon \}}$$

**Figure 7.** DScaleness Module Typing Rules

and here is the specification of lifting.

$$\frac{ST(p) = \bar{1} \quad \text{fields}(C) = \bar{T} \bar{f} \quad \bar{1} \xrightarrow{\text{lift}} \bar{d}, \bar{e} \quad x \text{ fresh}}{(p, C(\bar{R})) \xrightarrow{\text{lift}} \bar{d} @ (\llbracket C(\bar{R}) \rrbracket) x = \{ \bar{f} = \bar{e} \}), x}$$

### 4.4 DScaleness Type Checking

The primary novelty of DScaleness are the rules for DnesT module typing and composition, and that is the focus of this section. We adopt the typing rules of FGJ in their entirety, and refer the reader to [17] for relevant details.

DScaleness syntax for expressing DnesT module types is  $T \circ \mu\tau$ , where  $\mu\tau$  is a DnesT module type. The  $T$  in this form represents the type bounds of dynamically constructed types that have been used to instantiate the module; we refer to this part of the type as the *instance coercion*. Because these types are dynamically constructed, their identity is not known statically, hence the need to treat them as upper-bounded type names in the static type analysis. This subtle technical point of our type system is discussed at more length in Sect. 5. It is important to note that the type names in  $T$  will be fully resolved at run time, so that any module generated by a DScaleness program execution will have a fully reified DnesT type. Throughout this paper, we abbreviate  $\emptyset \circ \mu\tau$  as  $\mu\tau$  for brevity.

This is reflected in the MODT rule in Fig. 7, which connects the DnesT typing system with the DScaleness type system. Since in this case we are typing an uninstantiated module definition its instance coercion is empty. An instance coercion in a module type is directly populated when a module is instantiated, as in the MODINSTT rule. Here, the type instances  $\bar{e}_1$  are all dynamically constructed, so they define the upper bounds of the instantiated module’s instance coercion. We also expect all type and value parameters to respect the typing bounds specified in the module definition. The MODWIRET typing rule for module wiring is a straightforward reflection of the operational rule for module wiring, as is the MODIMAGET rule for module runnability imaging.

EXAMPLE 4.3. *Returning to the code and type examples in Sect. 2, we may assign the following typing:*

$$G \vdash \text{authSpecialize}(50, \text{radioM}, [[k_1, k_2]]) : \text{commT}$$

Given `radioM : radioT, k1 : uint8[], k2 : uint8[]` ∈  $G$ .

## 5. Scalness/nesT Foundations

The Scalness/nesT type system design is based on principles studied in the foundational calculus  $\langle \text{ML} \rangle^2$  [20].  $\langle \text{ML} \rangle$  comprises  $F_{\leq}$ , state, dynamic type construction, and staging features. In this section we describe how the design of modules and module operations in Scalness can be modeled in  $\langle \text{ML} \rangle$ . Although the correspondence is informal, these models directed the design of Scalness semantics and type checking, and provide confidence in its soundness. While our choice of modules as the basic unit of nesT code is based on obvious software engineering concerns and the need for a tight relation with nesC, Scalness modules are well correlated with certain structures in  $\langle \text{ML} \rangle$  and so are also technically appealing.

**The model of a module.** Code as a datatype is available in  $\langle \text{ML} \rangle$  as expressions of the form  $\langle e \rangle$ . While code as a datatype is a standard feature of staged/generative programming,  $\langle \text{ML} \rangle$  has adapted staged programming to a setting where different code levels are intended for execution on different machines with distinct process spaces. In particular, values, including code values, must be closed. If a type or term variable occurs free in  $\langle e \rangle$ , it must be  $\Lambda$  or  $\lambda$  bound, respectively, for closure. Hence, if a type variable  $t$  is free in  $\langle e \rangle$ , then  $\Lambda t \preceq \tau.\langle e \rangle$  binds it, and provides parametric subtyping polymorphism for  $\langle \text{ML} \rangle$  terms.

If the term variable  $x$  is free in  $\langle e \rangle$ , then  $\lambda x : \tau.\langle e \rangle$  binds it. Furthermore, the type  $\tau$  in the term  $\lambda x : \tau.\langle e \rangle$  must be of the form  $\langle \varsigma \rangle$ , because the type discipline requires that  $x$  is of code type, since it occurs within code. If the programmer wishes to pass a value residing at the current execution stage to such a function, it must be explicitly “lifted” in the now-standard style of [33]. However, in  $\langle \text{ML} \rangle$ , lifting a value entails serialization of it, which is non-trivial in case the value is stateful.

We use  $\langle \text{ML} \rangle$  type and term bindings to model Scalness type and term parameters. This is a standard strategy, in fact FGJ typing [17] is based on it as well. Hence the basic analog of a module is:

$$\Lambda t \preceq \tau.\lambda x : \langle \varsigma \rangle.\langle e \rangle$$

where  $t$  is a bounded type parameter and  $x$  is a value parameter.

**The model of instantiation.** Most of the interesting parts of Scalness typing happen at instantiation. Given the above model of a module, the  $\langle \text{ML} \rangle$  analog of instantiation is a term of the form:

$$(\Lambda t \preceq \tau.\lambda x : \langle \varsigma \rangle.\langle e \rangle)(\tau')(\text{lift } v)$$

where all parameters are instantiated. Note in particular that the value parameter  $v$  must be explicitly lifted, since the model must reflect that values passed in to modules are always constructed at the first stage in a Scalness program. This means that  $v$  must be assumed to not be a code value, while the type annotation on  $x$  requires that it be lifted. There is no explicit lift operation in Scalness, but the DScalness semantics (Fig. 6) specifies that serialization is always implicit at module instantiation. Scalness typing of instantiation thus treats value instantiation as  $\lambda$  application with implicit lifting of the argument, and type instantiation as  $\Lambda$  application, i.e. a form of bounded  $\forall$ -elimination.

**Type construction and variable escape.** A central technical novelty and core feature of DScalness is dynamic type construction for module instantiation. As we discussed in Sect. 2, this feature is technically challenging since constructed types can escape their scope of declaration. Similarly, in  $\langle \text{ML} \rangle$ , types may be dynamically constructed that can escape their declaration scope, in particular if they are used as function type annotations. An  $\exists$  type binder was introduced in  $\langle \text{ML} \rangle$  for this purpose; intuitively a type of the form  $\exists t \preceq \tau.\varsigma$  is a type containing a dynamically constructed type term  $t$  with upper bound  $\tau$ .  $\langle \text{ML} \rangle$  includes a “tlet” expression form for

constructing types, so for example:

$$\begin{aligned} \text{tlet } t \preceq \mathbf{uint16} = \text{if } e \text{ then } \mathbf{uint8} \text{ else } \mathbf{uint16} \text{ in } (\lambda x : t.x) : \\ \exists t \preceq \mathbf{uint16}.t \rightarrow t \end{aligned}$$

Here a type  $t$  is dynamically constructed to be either  $\mathbf{uint8}$  or  $\mathbf{uint16}$ , and then used in the type annotation of a type-specialized identity function. Furthermore,  $t$  escapes its declaration scope since it annotates a function argument. Since  $e$  is some arbitrary computation, we cannot statically predict what  $t$  will be, other than “some type with upper bound  $\mathbf{uint16}$ ”. Note also that since  $t$  can appear in contravariant positions, it is unsound to perform a covariant substitution of  $\mathbf{uint16}$  for  $t$ , so the  $\exists$  bound is needed. Although this usage of  $\exists$  types is somewhat non-standard, an eigenvariable interpretation of the bound type variable is sound and also consistent with standard existential type interpretations.

Inspired by these foundations, in DScalness the type form:

$$T_1 \circ \langle T_2; V \rangle \{ \iota; \varepsilon \}$$

captures the same typing mechanisms, in particular the instance coercion  $T_1$  is the analog of  $\exists$  bound type variables, in contrast to the type parameters  $T_2$  which are implicitly  $\forall$  bound, as discussed above. The static semantics of  $T_1$  and  $T_2$  are distinguished appropriately, especially in the treatment of the typing rules for module instantiation and module wiring in Fig. 7.

## 6. Implementation

Scalness is implemented as a modification to the open source Scala compiler. Scala was chosen as a basis for our first stage language because the Scala compiler supported a plug-in architecture, and we originally envisioned implementing Scalness as a plug-in. We also wanted to create a practical programming system, and Scala’s easy access to Java libraries and broad community support were attractive.

Keeping the Scala and Scalness type checking well separated in the implementation had useful software engineering benefits as well. It simplifies the problem of tracking the evolving Scala compiler. It also promotes a clear separation of first and second stage concerns in the mind of the Scalness programmer. Unfortunately we found that the needed modifications to the type checker could only be made by direct modification of the compiler code base.

In addition to type checking, runtime support is needed to implement Scalness module operations. Also, facilities are required to read nesT modules from the file system and parse them into ASTs, and to write TinyOS image source code files defined by constructed nesT modules at image invocations.

### 6.1 Online Repository and Examples

The Scalness/nesT compiler and several code examples, including applications discussed in Sect. 2 and Sect. 7, are available for download from the following URL:

<http://tinyurl.com/a85z8cu>

### 6.2 nesT Type Checking and Program Transformation

The nesT language is treated by two major components in the implementation, the type checker and the nesT-to-nesC rewriting transformation. The nesT type checker was written from the ground up, in contrast to the Scalness type checker which was defined as an extension to the Scala type checker. The rewriting transformation yields TinyOS2-compliant source code, which can be separately compiled.

The nesT language is defined as a subset of the nesC language. An AST yielded by parsing is type checked by our algorithm, which incorporates subtyping and other features not present in nesC type

<sup>2</sup>Pronounced “framed ML.”

checking. This algorithm is a nearly direct encoding of the type discipline described in Sect. 3. Following type checking, the AST is submitted to a rewriting transformation that imposes semantic disciplines discussed in Sect. 3, in particular type safe casting and array bounds checks, also in nesC. For example, a statement of the form `x = a[e]` will be rewritten to:

```
int _x = e; if (_x >= a_SIZE) fail(); x = a[_x];
```

where `a_SIZE` is an automatically generated variable containing the size of `a` and `fail` is some user-defined function that handles array bounds check failure.

Source code for nesT module definitions is written in separate files that are included in Scalanness code, as discussed below. This separation is mainly for software engineering purposes, to avoid modifying the Scala compiler to parse intermingled Scala and nesT syntax.

### 6.3 Scalanness Module Language Syntax

In order to limit modifications of the Scala compiler and reduce engineering problems in our implementation, we have avoided modifying Scala syntax to represent Scalanness features. Hence, modules are represented as class instances, which must satisfy the following trait:

```
trait NesTModule {
  def image(): Unit // Generates residual nesC program.
  def +>(m: NesTModule): NesTModule // Wires this to m.
}
```

This trait is implemented by a `NesTModule` class that provides the appropriate semantics for wiring and TinyOS image generation, including translation to nesC and file output. This class also manages parsing and storage of nesT ASTs from source code files, and type checking of nesT ASTs.

Any nesT module definition is a subclass of `NesTModule`. Some subtleties are involved in supporting first class *generic* modules. Instantiation is implemented by method call, but since type and value parameters vary per module, particular modules must define their own parameters and instantiation methods. For example, we would represent the `authSend` component definition from Sect. 2, line 9 as follows:

```
class authSend extends NesTModule {
  var mt : MetaType[LiftableType] = _
  var sendk : LiftableType = _
  def instantiate(m: MetaType[LiftableType], k: LiftableType) =
    { val result=new authSend; result.mt=m; result.sendk=k }
  "authSend.nt"
}
```

Although the `instantiate` method and parameter fields must be defined in the implementation at the time of this writing, compiler generation of these definitions is a topic for future work; any modules `instantiate` method can be easily inferred from its type annotation. Note that the types at which parameters are declared are as general as possible (e.g. `s` and `n` are not declared as `uints` but as `LiftableTypes`). This is because class definitions support the semantics of Scalanness, not Scalanness type checking (discussed below), and declaring generic parameters at a maximally general type removes interference related to Scala type checking. Finally, note the string literal `"authSend.nt"` at the end of the definition. This is the file containing the nesT source code definition of the module. The Scala compiler has been modified to input and parse the specified source code when this literal is encountered during the Scala type checking phase.

### 6.4 Type Annotation and Checking

Scalanness typing relies on native Scala syntax for terms, specifically Scala annotations and singleton types are utilized. Scala annotations allow metadata to be associated with definitions. A mod-

ule type annotation is of the form `@ModuleType("μr")`, where `μr` is defined using the syntax of Fig. 3. The compiler-defined `ModuleType` class automatically associates the type with the identifier immediately following it. In the case of module class definitions, the type is assigned as a class field. In the case of variable definitions, the type is assigned as a Scala singleton type of the object. For example, the declaration of `authSend` on line 9 in as in Sect. 2 would be preceded by such an annotation where `μr` is the type specified in Example 3.1, and `sendM` as on line 22 would be annotated with an instance of that type. Similarly, annotations are required on method parameter and result types, if those methods expect nesT modules as arguments or return them, as for the `radioC` parameter of the `authSpecialize` method defined in Sect. 2, and the method's `commT` return type. These requirements reflect the type discipline in Scalanness as specified in Sect. 4, which requires module type annotations at these points.

Scalanness type checking has been implemented as an analysis of these annotations during Scala type checking, piggybacking on that process. When type checking a class that extends `NesTModule`, the compiler uses its type annotation to perform nesT type checking on the underlying AST representation of the module. When type checking module operations (i.e. at invocations of `instantiate`, `+>`, or `image`), the Scala compiler has been modified to examine operand types for Scalanness type annotations, and to decorate resultant singleton types of these operations with new Scalanness annotations, reflecting the typing rules in Fig. 7. A type checking exception is raised in case this analysis fails. Scalanness type checking does not modify Scala type checking in any other way, so it is a conservative extension of Scala typing.

### 6.5 Importing nesC Libraries

Our preliminary experiments with nesT show that it is expressive enough to write useful program components. However, any realistic application will need to interact with various libraries written in nesC. One library of critical importance is the TinyOS operating system itself. Our current solution is to allow non-generic nesC components to be treated as nesT modules as long as they only use or provide commands, which are interpreted as nesT imports and exports. Support for specializable generic nesC library components is a topic for future work. Events can be accessed through "shim" modules provided by the user, since used or provided events are really just syntactic sugar for provided or used commands respectively. A library component defined in a file `LibraryC.nc` can be defined as a nesT module as follows:

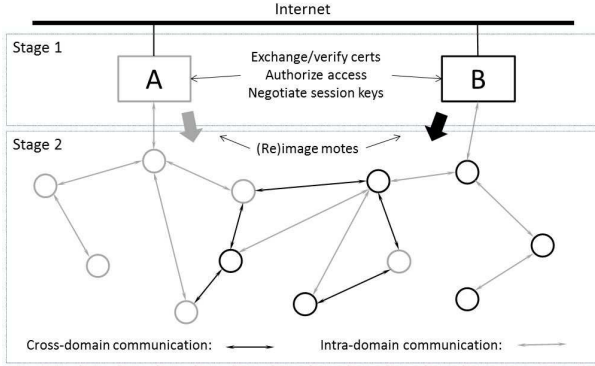
```
object LibraryC extends NesTModule { external("LibraryC.nc") }
```

Note that nesC code imported in this way is not type checked by the Scalanness/nesT compiler, since nesT is a strict subset of nesC. Rather, the programmer type annotates the shimmed module using a `@ModuleType` annotation as for other module definitions, and the compiler trusts that the annotation is correct. This introduces a possibility for type safety failure in our system, if the imported code contains a type error. A possible long term goal would be a complete re-write of TinyOS in nesT, yielding full type safety of all sensor code, but this is well beyond our current scope.

## 7. Example: Authorization and Access Control

In [5] the SpartanRPC architecture for link-layer resource authorization in TinyOS-based WSNs is developed (as an extension of [4]). In SpartanRPC, resources are accessed by link-layer remote procedure calls (RPC) which require authorization. Users are authorized by communicating credentials to the provider, expressed in an authorization logic based on RT [18] and implemented using TinyECC [19] public key signatures. SpartanRPC supports an "open world" security model, allowing WSNs in different security





**Figure 8.** Staging Authorization and Authorized Access in a Multi-Domain WSN.

domains to interact without sharing secrets *a priori*. However, public key encryption and signature verification is very expensive in a WSN. Hence, session keys are negotiated for ongoing resource access (using a TinyECC-based Diffie-Hellman protocol).

In this section we describe a re-implementation of the SpartanRPC protocol in Scalness/nesT that addresses several shortcomings of SpartanRPC, and will thus serve to illustrate the power of Scalness/nesT. The central idea, illustrated in Fig. 8, is that responsibility for authorization on the basis of public key credentials is offloaded from the WSN to a Scalness program running on a hub device or lab computer. We assume a WSN comprising two sub-networks under control of distinct security domains *A* and *B*. Each domain also controls a lab or hub device which is in communication with WSN nodes in their domain, either prior to or during deployment. These devices are in communication with each other over the Internet, and exchange authorization credentials for their domain over that medium in the first-stage Scalness program. Each device then confirms authorization for resource access according to their own domain’s policies, and subsequently they negotiate session keys over the Internet. These keys are then used to specialize nesT code for imaging on WSN nodes. The overall architecture of this application represents a concrete realization of the ideas of Fig. 1, and also expands on and implements the idealized example presented in Sect. 2.

Note that our current implementation assumes nodes are programmed in the lab since we have no secure OTA program dissemination library; the Deluge protocol has such an extension [9] that we plan on using in the future.

**Evaluation on Snowcloud** To empirically evaluate the staged implementation of SpartanRPC in Scalness/nesT, we have implemented and tested both the original SpartanRPC as well as the Scalness/nesT staged version in our deployed Snowcloud WSN system architecture. Mobile gateway devices as pictured in Fig. 2 are provided to Snowcloud system users for data gathering, and are also used by system administrators for controlling sampling rates. The hardware for both of these so-called “harvester” devices, the same for users and administrators, is equipped with a mote for establishing network communication. When the device is introduced to the sensor network, the two together comprise a single network with two distinct security domains – the sensor node subnetwork, and the subnetwork of the single device mote. The mote on harvester devices provided to system users is supplied with credentials for collecting data, but not modifying network control, whereas system administration harvester motes are supplied with stronger credentials for both functions.

The original and Scalness/nesT versions of this application can be compared both in terms of performance and user experience. In the unstaged version, the SpartanRPC protocol requires an initial network configuration period when credentials are exchanged and verified. Since a single TinyECC signature requires at least 90 seconds to verify on the Crossbow TelosB platform [5] with a fully dedicated processor, there is an initial network “warmup” period of at least a few minutes. Also, in the unstaged version, upon first invocation of an RPC service Diffie-Hellman is used in the network to negotiate a session key. In the staged version, credential exchange, validation, and session key negotiation are all performed on the high-powered hub. For this reason, mote code size in ROM is significantly reduced. There are differences in RAM usage as well, due to authorization overhead in the unstaged version and also the storage of key material in RAM vs. ROM, since specialization of code with key material in the staged version allows the latter. Note that this difference is intensified by scale and the number of keys (i.e. RPC services) needed by an application. Lower RAM and ROM usage can have significant performance impacts on deployed code. In the following table we summarize RAM and ROM usage for the harvester and sensor node images for three software versions: one with no security mechanisms in place, one with unstaged SpartanRPC protocols in place, and one generated by Scalness evaluation in our staged version of the SpartanRPC protocol.

		Unsecured	Unstaged	Staged	Savings
Sensor:	ROM	36254	48616	36596	25%
	RAM	2868	5417	3038	44%
Harvester:	ROM	24316	35834	24436	32%
	RAM	2274	4771	2402	50%

The “Savings” are the percent reduction from unstaged to staged secure implementation, and these numbers show the potential for saving both RAM and ROM space is significant. From the perspective of user experience, the staged version of this application is more convenient, since no initial authorization period is needed when the harvester is first introduced to the network. The staged version also exposes the system to fewer bugs and failures that would be obstacles to the primary goal of data collection.

## 8. Conclusion

We have introduced Scalness/nesT, a two stage programming system for wireless sensor networks. Our system provides a powerful programming environment for dynamically specializing and composing nesC modules in a type safe way; any type correct Scalness program will generate only type correct residual programs.

### 8.1 Related Work

We do not review the broader topic of sensor network programming here; the reader is referred to [28] for a broader perspective.

We follow the foundational (ML) work in our language design [20]; Sect. 5 discusses how it serves as the theoretical underpinning of our approach. The primary aim of this work is to make the theoretical insights of (ML) more practical. We accomplish this by making a sensor language nesT that is based on the design of nesC, and by implementing Scalness and nesT and testing the framework on examples.

The potential of applying metaprogramming to sensor networks was explored in the functional sensor language Flask [24]. Flask allows FRP-based stream combinators to be pre-computed before network deployment, but it is possible to generate ill-typed Flask object code since cross-stage static type checking is not performed. Hume [15] is a DSL for real-time embedded device programming.

It includes a metaprogramming layer but that layer is more like nesC's configuration files in that there is a very restricted syntax for a few special metaprogramming operations including component wiring, macros, and code templating.

MetaML [32, 33] and MetaHaskell [23] are a foundations we build on; they do not address type specialization or dynamic type construction.

Lightweight Modular Staging [31] describes a method of expressing staged computations using a Scala host framework without any compiler modifications. The approach allows cross-stage type safety but does not support dynamic type construction.

Actor based sensor metaprogramming has been studied in [6]; this work shares our focus on high level dynamic reprogrammability but is untyped. More broadly, meta programming is known to be useful for increasing the efficiency of systems applications. One example is Tempo [8], a system that integrates partial evaluation and type specialization for increasing efficiency of systems applications. Ur [7] allows for type safe meta programming for web applications.

The units of staged code composition in nesT programming are *modules*. Countless different module systems exist, but they are primarily designed to achieve separate compilation and sound linking [2]. Our different design goal leads to different design choices in nesT modules. For example, data crossing nesT module boundaries needs to conform to the property of process separation, a non-issue in standard module system designs. In addition, nesT modules allow values/types across the boundary of modules to be flexibly constructed, including dynamic construction of types, to achieve maximal flexibility of cross-stage specialization. Module systems such as ML modules [21] and Units [10] allow types to be imported/exported as we also support; there are several features of ML modules including type hiding that we do not aim to support. nesT modules are more expressive in their support of first class modules as values and the possibility of dynamic construction of "type exports." That said, first class modules are not new [1, 25], we only claim novelty in their application to program staging and the incorporation of dynamic type construction.

The type parametricity of System F and  $F_{\leq}$  [3], and the practical type systems it inspired such as Java's generics, do not treat types as first class values as we do. C++ templates support types as meta values in template expansion, but type safety of generated code is not guaranteed without full template expansion. Concepts [14] improves on this, but types are still not first class values.

## Acknowledgments

We acknowledge Yu David Liu for early contributions to this work.

## References

- [1] D. Ancona and E. Zucca. A calculus of module systems. *Journal of functional programming*, 11:91–132, 2002.
- [2] L. Cardelli. Program fragments, linking, and modularization. In *POPL*, pages 266–277, 1997.
- [3] L. Cardelli and P. Wegner. On understanding types, data abstraction, and polymorphism. *ACM Comput. Surv.*, 17(4):471–523, 1985.
- [4] P. Chapin and C. Skalka. SpartanRPC: Secure WSN middleware for cooperating domains. In *MASS*, November 2010.
- [5] P. Chapin and C. Skalka. Spartan RPC. Technical report, University of Vermont, 2013. Submitted. <http://www.cs.uvm.edu/~skalka/skalka-pubs/chapin-skalka-spartanrpc.pdf>.
- [6] E. Cheong. *Actor-Oriented Programming for Wireless Sensor Networks*. PhD thesis, University of California, Berkeley, 2007.
- [7] A. Chlipala. Ur: Statically-typed metaprogramming with type-level record computation. In *PLDI*, 2010.
- [8] C. Consel, L. Hornof, R. Marlet, G. Muller, S. Thibault, E.-N. Volanschi, J. Lawall, and J. Noyé. Tempo: specializing systems applications and beyond. *ACM Comput. Surv.*, 1998.
- [9] P. K. Dutta, J. W. Hui, D. C. Chu, and D. E. Culler. Securing the deluge network programming system. In *IPSN*, pages 326–333, 2006.
- [10] M. Flatt and M. Felleisen. Units: Cool modules for HOT languages. In *PLDI*, 1998.
- [11] J. Frolik and C. Skalka. Snowcloud: A complete system for snow hydrology research. In *RealWSN*, 2013.
- [12] D. Gay, P. Levis, R. von Behren, M. Welsh, E. Brewer, and D. Culler. The nesC language: A holistic approach to networked embedded systems. In *PLDI*, 2003.
- [13] G. Ghelli and B. Pierce. Bounded existentials and minimal typing. *Theoretical Computer Science*, 193(1-2):75 – 96, 1998.
- [14] D. Gregor, J. Järvi, J. G. Siek, G. D. Reis, B. Stroustrup, and A. Lumsdaine. Concepts: Linguistic support for generic programming in C++. In *OOPSLA*, 2006.
- [15] K. Hammond and G. Michaelson. Hume: A domain-specific language for real-time embedded systems. In *GPCE*, pages 37–56. Springer-Verlag, 2003.
- [16] J. Hill, R. Szewczyk, A. Woo, S. Hollar, D. E. Culler, and K. S. J. Pister. System architecture directions for networked sensors. In *ASPLAS*, pages 93–104, 2000.
- [17] A. Igarashi, B. C. Pierce, and P. Wadler. Featherweight Java. *ACM Trans. Program. Lang. Syst.*, 23(3):396–450, 2001.
- [18] N. Li and J. C. Mitchell. RT: A role-based trust-management framework. In *Proceedings of the 3rd DARPA Information Survivability Conference and Exposition*, pages 201–212, 2003.
- [19] A. Liu and P. Ning. Tinyecc: A configurable library for elliptic curve cryptography in wireless sensor networks. In *IPSN*, pages 245–256, 2008.
- [20] Y. Liu, C. Skalka, and S. Smith. Type-specialized staged programming with process separation. *HOSC*, pages 341–385, 2011.
- [21] D. MacQueen. Modules for Standard ML. In *Proceedings of ACM Conference on Lisp and Functional Programming*, 1984.
- [22] S. Madden, M. J. Franklin, J. M. Hellerstein, and W. Hong. TAG: a Tiny AGgregation service for ad-hoc sensor networks. *SIGOPS Oper. Syst. Rev.*, 36(SI):131–146, 2002.
- [23] G. Mainland. Explicitly heterogeneous metaprogramming with MetaHaskell. In *ICFP*, 2012.
- [24] G. Mainland, G. Morrisett, and M. Welsh. Flask: staged functional programming for sensor networks. In *ICFP*, 2008.
- [25] J. Mitchell, S. Meldal, and N. Madhav. An extension of standard ML modules with subtyping and inheritance. In *POPL*, 1991.
- [26] C. D. Moeser, M. Walker, C. Skalka, and J. Frolik. Application of a wireless sensor network for distributed snow water equivalence estimation. In *Western Snow Conference*, 2011.
- [27] T. Molhave and L. H. Petersen. Assignment Featherweight Java. Master's thesis, University of Aarhus, 2005.
- [28] L. Mottola and G. P. Picco. Programming wireless sensor networks. *ACM Computing Surveys*, 2011.
- [29] G. C. Necula, S. McPeak, and W. Weimer. CCured: type-safe retrofitting of legacy code. In *POPL*, 2002.
- [30] M. Odersky, L. Spoon, and B. Venners. *Programming in Scala, second edition*. Artima, Inc, 2011.
- [31] T. Rompf and M. Odersky. Lightweight modular staging: a pragmatic approach to runtime code generation and compiled DSLs. In *GPCE*, pages 127–136, 2010.
- [32] W. Taha. Resource-aware programming. In *ICISS*, pages 38–43, 2004.
- [33] W. Taha and T. Sheard. Multi-stage programming with explicit annotations. In *PEPM*, pages 203–217, 1997. ISBN 0-89791-917-3.
- [34] R. Willett, A. Martin, and R. Nowak. Backcasting: adaptive sampling for sensor networks. In *IPSN*, pages 124–133, 2004.