

A TYPE SYSTEM FOR SAFE SN RESOURCE ALLOCATION

Michael Ocean Assaf Kfoury Azer Bestavros
Computer Science Department
Boston University
Boston, MA 02215

Technical Report: BUCS-TR-2008-011

June 14, 2008

Abstract

SNBENCH is a platform on which novice users compose and deploy distributed Sense and Respond programs for simultaneous execution on a shared, distributed infrastructure. It is a natural imperative that we have the ability to (1) verify the safety/correctness of newly submitted tasks and (2) derive the resource requirements for these tasks such that correct allocation may occur. To achieve these goals we have established a *multi-dimensional* sized type system for our functional-style Domain Specific Language (DSL) called Sensor Task Execution Plan (STEP). In such a type system data types are annotated with a vector of size attributes (*e.g.*, upper and lower size bounds). Tracking multiple size aspects proves essential in a system in which Images are manipulated as a first class data type, as image manipulation functions may have specific minimum and/or maximum resolution restrictions on the input they can correctly process.

Through static analysis of STEP instances we not only verify basic type safety and establish upper computational resource bounds (*i.e.*, time and space), but we also derive and solve data and resource sizing constraints (*e.g.*, Image resolution, camera capabilities) from the implicit constraints embedded in program instances. In fact, the static methods presented here have benefit beyond their application to Image data, and may be extended to other data types that require tracking multiple dimensions (*e.g.*, image “quality”, video frame-rate or aspect ratio, audio sampling rate). In this paper we present the syntax and semantics of our functional language, our type system that builds costs and resource/data constraints, and (through both formalism and specific details of our implementation) provide concrete examples of how the constraints and sizing information are used in practice.

[†] This research was supported in part by a number of NSF awards, including CISE/CSR Award #0720604, ENG/EFRI Award #0735974, CISE/CNS Award #0524477, CNS/NeTS Award #0520166, and CNS/ITR Award #0205294.

1 Introduction

Motivation

The Sensor Network WorkBench (SNBENCH) is a collection of compile-time tools and run-time components that enable the painless development and deployment of Sense and Response services that run on a shared infrastructure. Toward SNBENCH’s goal of enabling novice users to compose these services we provide our users with a functional-style Domain Specific Language (DSL) for specification, called STEP (Sensor Task Execution Plan).¹ STEP is resource agnostic insofar as service logic may refer to particular types of resources (*e.g.*, an Image sensor) without indicating which *specific* resources should be utilized within the service.

Our ability to allocate resources on which to deploy STEP services is contingent upon our ability to verify the safety of new services and to derive resource requirements from new service instances. While the motivation of the SNBENCH project and specific implementation details for its architecture and run-time components have been published elsewhere ([OBK06]), in this paper we present the static analysis techniques that we have developed to provide safety and resource constraint extraction on our sensing-centric STEP language. We base our type system on sized (*a.k.a.* static-dependent) type systems, wherein upper bound size annotations on types coupled with cost functions are used to determine memory (storage) and processing (worst case execution) bounds.

We expand our size tracking to multiple dimensions (*i.e.*, multiple dimensions of size annotations) toward the goals of (1) supporting Images as a first-class data type and (2) enabling the static inference of required image (and image sensor) resolutions from implicit constraints.

Unlike traditional scalar data, both size bounds of an Image (*i.e.*, upper and lower, where the lower bound is the potential minimum image resolution) may have an impact on functional correctness. For example, attempting to recognize a face in a low-resolution image may never succeed, or worse, might diverge depending on the implementation. While one could consider adding additional types and subtyping relations to the type system to support awareness of image resolutions (*e.g.*, LowResolutionImage, MediumResolutionImage, HighResolutionImage) it should be obvious that this sort of solution does not scale.

Our sized type system can not only bound costs for memory and computation, but also produces sensing domain specific resource constraints. In tracking bound both upper size bounds and lower size bounds we are able to make statements that bound a worst-case execution time and also provide bounds for Image resolution; the latter property ultimately leads to establishing the correctness of Image processing expressions.

A Motivating Example

Our goal is to be able to leverage the size annotations in the type system to provide both an upper-bound for computational requirements of services (as prior works have done), while additionally (1) maintaining

¹We actually provide other, high-level languages that are compiled down to STEP as our common Instruction Set Architecture.

minimum size aspects to verify correctness in the presence of functions that require a minimum size to ensure correctness and (2) determining and maintaining implicit constraints on resources and data sizes as extracted from contextual usage in a given service instance.

For example consider the code fragment below:

```
(* every 100 milliseconds, take an image from "any" camera and
try to detect motion. If motion is detected, send an e-mail. *)
letonce img = get(sensor(IMAGE,ANY)) in
    period(100ms,
        trigger(facedetect(img),email("mocean",img))
    )
```

In the code given, the variable `img` represents an image captured from “any” image sensor. However not all image sensors (*i.e.*, cameras) have the same capabilities with respect to image resolution (*e.g.*, a webcam might capture images at a resolution several times lower than that of an embedded Pan-Tilt-Zoom camera). In this program instance there are *implicit* constraints on the image that indicate that not just *any* sensor will do. The function (or as we call it, Opcode) `facedetect` constrains the size of `img`, as it requires a minimum resolution to correctly detect a face in an image. While the explicit periodicity function (or as well call it, Flowtype) `period` indicates that this expression must run every 100 milliseconds. Thus there is another constraint on the resolution; the resolution must also be low enough to allow computation every 100 milliseconds. These constraints on the resolution of the image must propagate back to the image sensor from which the image will be acquired to ensure that the sensor reserved for this program can support the required resolution (or range of resolutions).

Our sized constraint set (when solved) can be used to (1) guide task assignment (*e.g.*, do not split computation over the network where the data size will incur a steep networking overhead), (2) guide resource allocation (*e.g.*, reserve the correct sensor determined from a resolution range derived from use in context), and (3) determine if a program is fundamentally or temporarily infeasible (*e.g.*, some specific resolution too low to perform computations, required periodicity can not be met given current available resources).

In this paper we present our type system as applied to a subset of our domain specific language, give details of its implementation, and work through concrete examples of the system in use. The organization of the rest of this paper is as follows: Section 2 provides a brief, high-level overview of SNBENCH, Section 3 provides the syntax of our DSL and the static and dynamic semantics of our type system, and Section 4 applies this formalism to some examples to show the system in action as well as giving an overview of our implementation. In Section 5 we indicate some of the many possible future directions for this work.

Related Work

Bounding the execution time of programs and program fragments is a well-established problem in computing. There was a large interest in applying custom type systems to domain specific languages in the late nineties (*e.g.*, the USENIX Conference on Domain-Specific Languages (DSL) in 1997 and 1999).

Indeed, Our work has been largely inspired by existing works that aim to solve execution time bounding via an upper bound size annotation on types within a type system. These works include Static Dependent Costs [RG94], Sized Type Systems [HPS96], and Sized Time Systems [LH96]. All establish a formal type system that has been annotated with upper size bounds on datatypes to estimate an upper bound on execution time and memory requirements given input size. We intended to use these techniques (with small adjustments to support STEP) in order to verify basic type safety and extract execution and memory bounds to guide resource allocation and scheduling within SNBENCH. The operating environment of SNBENCH is intrinsically distributed and time dependent, yet it seemed that these existing works should have been able to be ported more or less directly ([HFH06], for example, is the current incarnation of the Sized Time System, and is aimed directly at the real-time embedded sensing community). However, this was not the case.

Our language (and infrastructure) supports the direct modification of Image data which, unlike traditional scalar data, has an overloaded notion of size (*i.e.*, resolution) that has a direct impact on functional correctness. In this environment the type signature of a function must include explicit resolution (size) bounds to convey what size *ranges* of data a function can *correctly* process. Thus our needs began to diverge almost immediately. In the existing works size annotation has nothing to do with functional correctness, moreover we recognized a need to track a lower size bound (annotation) in addition to the upper size bound. From the need to add the additional lower size bound we have established a system in which the size annotations are *multi-dimensional*; while the formalism described in section 3 only includes a lower and upper size bound, Section 5 discusses how easily more dimensions could be added and gives examples.

Additionally our system uses size constraints to solve for data size when it is not explicitly specified by the programmer (*i.e.*, size annotation variables). Our constraint set is explicit within the typing rules yet constraints are derived implicitly from program code, using our constrained size type signature for primitive operators. In solving the constraint set we can deduce feasible (and/or optimal) data sizes which directly map to image resolutions, resource constraints and sensor capabilities for image manipulation programs. Finally we allow primitive operators that directly manipulate the constraint set, allowing the programmer to explicitly constrain types without influencing the execution (so called, Flowtypes). We are unaware of any other such work that treats images as first class datatypes or that uses a type system to statically create such size constraint relationships to deduce required image sizes and sensor capabilities.

It is also worth noting that our language includes a slightly unusual let semantic which is advantageous in time dependent programming environments. While several functional languages include this same let behavior in their implementation, in time-independent operating environments implementation of the let discipline in this way does not influence the result of the computation, but rather is provided as an optimization. In our environment, this is not the case. Thus, we define our let semantic formally and include it in our proofs for soundness and completeness.

2 snBench Overview

To orient the reader to the platform to ease further discussion, in this section we briefly highlight the salient features of SNBENCH. The vision, goals and high-level overview of the SNBENCH infrastructure have been reported elsewhere [BBKO05] and implementation details may be found in [OBK06].

SNBENCH consists of programming support and a runtime infrastructure for Sensor Networks comprised of heterogeneous sensing and computing elements that are physically embedded into a shared environment. We refer to such a physical space with an embedded SN as a Sensorium [Bos]. The SNBENCH framework allows Sensorium users to easily program, deploy, and monitor the services that run in this space while insulating the user from the complexity of the physical resources therein. We liken the support that SNBENCH extends to a Sensor Network to the support that higher-level languages and operating systems provide to traditional, single machine environments (language safety, APIs, virtualization of resources, scheduling, resource management, *etc.*). SNBENCH is extensible by design such that new hardware and software capabilities may be painlessly folded into the infrastructure by its advanced users and those new capabilities easily leveraged by its novice users.

SNBENCH provides a high-level programming language with which to specify programs (services) that are submitted to the resource management component which in turn disseminates program fragments to the run-time infrastructure for execution. At the lowest level, each sensing and/or computing element hosts a Sensor eXecution Environment (SXE) that abstracts away specific details of the host and attached sensory hardware. SXEs are assigned tasks by the resource management components of SNBENCH the Sensorium Service Dispatcher and Sensorium Resource Manager in tandem monitor SN resources, schedule (link) and deploy (bind) tasks on to available SXEs. A graphical representation of this end-to-end support is shown in Figure 1.




Figure 1: The SN program life-cycle as enabled by snBENCH. Rectangles represent data, circles represent tasks/processes, and the dashed lines represent control communication (dependency).

The Virtual Instruction Set Architecture of snBENCH is the Sensorium Task Execution Plan (STEP), a domain specific tasking-language used to describe complete programs and fragments alike. A STEP program is a graph of a SN programs's data-flow and computational dependency with the nodes of a STEP graph representing the atomic computation and sensing operations and edges representing data flow. In execution, demand for evaluation pushed down from the root of the graph to the leaves and values percolate up from the leaves back to the root. STEP nodes describe data, control flow (*e.g.*, repetition, branching) and computation operations that we refer to as STEP Opcodes and the SXE

maintains implementations of the Opcodes with which it may be tasked.

STEP tasks (or subtasks) are assigned (linked) to the available SXEs for execution by the Service Dispatcher. The Service Dispatcher must, therefore, be able to characterize the STEP tasks and ensure that the resources required by the tasks are currently available in the system. We make a distinction between uncertified STEP tasks and those that have been certified (*i.e.*, validated by our type system as correct and annotated with resource requirements).² It is our type system that provides this certification and the resource requirements it extracts are essential to enabling the Service Dispatcher to correctly allocate resources (SXEs).

3 Formalism for “Core” STEP

In this section we present the formal logic that underlies the verification component described above. The type system we present in this Section supports a subset of the complete STEP programming language; we call this subset “Core” STEP. We discuss the challenges of supporting the few remaining STEP components in Section 5.

Readers who are actively familiar with the project may note that what is presented as STEP appears to be STEP’s high-level, functional sibling SNAFU. In fact, SNAFU is a largely a convenience wrapper for STEP, including some additional syntactic sugar (which we do not address in this paper).

In Core STEP we formally present only one let form. As we show later, this let behaves in a way that is temporally interesting. The Complete STEP supports other, traditional let forms (*i.e.*, lazy and eager), yet we omit them from the Core formalism as they are well-known and relatively less interesting in our domain.

3.1 Syntax of Expressions

Below we define the syntax of the valid expression forms of Core STEP.

²Certification could happen entirely on the client side in an environment in which the user/service composer can not forge certification output.

$e ::=$		expressions
	v	<i>value</i>
	x	<i>variable</i>
	$\text{cond } e \ e \ e$	<i>conditional (if-then-else)</i>
	$\text{let } \{x_i=e_i\}^{i \in 1..n} \text{ in } e$	<i>let binding</i>
	$\text{get } e$	<i>read a sensor</i>
	$op \ e$	<i>opcode/primitive operation</i>
	$\text{trigger } \{x_i=e_i\}^{i \in 1..n} \ e \ e$	<i>construct for repetition</i>
$v ::=$		values
	$0 \mid 1 \mid 2 \mid \dots$	<i>integer</i>
	$\text{true} \mid \text{false}$	<i>boolean</i>
	i	<i>image</i>
	$\text{time} \mid \text{image}$	<i>sensor</i>
$op ::=$	$op_1 \mid op_2 \mid op_3 \mid \dots$	opcodes

3.2 Preliminary Definitions

The evaluation (dynamic semantics) of expressions take the general form:

$$e \mid \nu \mid t \rightarrow e' \mid \nu' \mid t + 1$$

Where ν is a special variable store required for our let semantic, and t is a discretized time that increments per evaluation step (*e.g.*, computation count, clock tick). We read this in English as “the expression e and its variable store ν take an evaluative step to the new expression e' with new variable store ν' .”

3.2.1 Definition: The Variable Store ν :

The evaluation of the STEP language is time dependent insofar as values that are read from the sensing environment may change over time. Thus, *when* a function is evaluated will directly effect the value retrieved (specifically *via* the sensor reading function, `get`). Put differently, the let semantic we provide has a direct impact on the values retrieved for a variable (*e.g.*, an eager let retrieves a reading at time 0, a typical lazy or by-need let retrieves a new reading at every time the substituted variable is encountered). To best accommodate the needs of STEP and this temporal dependence, we define our default let-binding to be a hybrid approach; we offer a deferred evaluation (*i.e.*, by need) that also offers the reuse provided by eager evaluation (*i.e.*, with caching).

To achieve the desired result our let-binding construct manipulates a variable store ν that stores the mapping of variables to unreduced expressions which will be reduced within the environment itself when the variable is encountered (by need). Should the same variable be encountered again further in the expression, the mapping in ν will later point to the fully reduced value and will therefore return that previously computed value (with caching).

We define the store, ν below, whose domain ranges over variable symbols, each pointing to valid STEP expression (either an unreduced expression or a value).

$$\begin{aligned}\nu &= \{x_1 \mapsto e_1, x_2 \mapsto e_2, \dots, x_n \mapsto e_n\} \\ \text{Domain}(\nu) &= \{x_1, \dots, x_n\} \\ \nu(x_i) &= e_i\end{aligned}$$

3.2.2 Definition: The free-variable function FV

FV calculates the free-variables in an expression (e) or expression and store pairing ($e \mid \nu$).

$$\begin{aligned}FV(v) &= \emptyset \\ FV(x) &= \{x\} \\ FV(\text{cond } e_1 \ e_2 \ e_3) &= FV(e_1) \cup FV(e_2) \cup FV(e_3) \\ FV(\text{let } \{x_i = e_i\}^{i \in 1..n} \text{ in } e) &= \left(\bigcup_{i=1}^n FV(e_i) \cup FV(e) \right) - \{x_1 \dots x_n\} \\ FV(\text{get } e) &= FV(e) \\ FV(\text{op } e) &= FV(e) \\ FV(\text{trigger } \{x_i = e_i\}^{i \in 1..n} \ e_{n+1} \ e_{n+2}) &= \left(\bigcup_{i=1}^n FV(e_i) \cup FV(e_{n+1}) \cup FV(e_{n+2}) \right) - \{x_1 \dots x_n\} \\ FV(e \mid \nu) &= \left(\bigcup \{FV(\nu(x)) \mid x \in \text{Domain}(\nu)\} \cup FV(e) \right) - \text{Domain}(\nu)\end{aligned}$$

3.2.3 Definition: expression and store pair closure

We say $e \mid \nu$ is **closed** if and only if $FV(e \mid \nu) = \emptyset$.

3.3 Dynamic Semantics

In this section we present the evaluation rules (dynamic semantics) for the various constructs (*i.e.*, syntactic forms) of the language.

3.3.1 Conditional

$$\begin{aligned}(\text{E-IFTRUE}) \quad \text{cond true } e_2 \ e_3 \mid \nu \mid t &\rightarrow e_2 \mid \nu \mid t + 1 \\ (\text{E-IFFALSE}) \quad \text{cond false } e_2 \ e_3 \mid \nu \mid t &\rightarrow e_3 \mid \nu \mid t + 1 \\ (\text{E-IF}) \quad \frac{e_1 \mid \nu \mid t \rightarrow e'_1 \mid \nu' \mid t + 1}{\text{cond } e_1 \ e_2 \ e_3 \mid \nu \mid t \rightarrow \text{cond } e'_1 \ e_2 \ e_3 \mid \nu' \mid t + 1}\end{aligned}$$

3.3.2 Opcodes

We refer to primitive operators in STEP as Opcodes. Evaluation of opcodes is strict (arguments must be reduced to values before the opcode may be evaluated). Evaluation and typing rules for specific instances of opcodes (*e.g.*, `facetect`, `resample`) are given in Section 4. The general form is presenter here.

$$(E-OP1) \frac{e_1 \mid \nu \mid t \rightarrow e'_1 \mid \nu' \mid t + 1}{op \ e_1 \mid \nu \mid t \rightarrow op \ e'_1 \mid \nu' \mid t + 1}$$

where $op \in \{op_1 \mid op_2 \mid op_3 \mid \dots\}$

$$(E-OPAPPLY) \quad op \ v_1 \mid \nu \mid t \rightarrow v_2 \mid \nu \mid t + 1$$

where $v_2 = \mathit{Apply}(op \ v_1)$
and $op \in \{op_1 \mid op_2 \mid op_3 \mid \dots\}$

3.3.3 Sensor reads (and the physical sensor environment \mathcal{E})

As STEP is a sensing centric DSL, it is essential that we have the ability to read values from sensors that are embedded in the physical environment. We imagine a logical array that contains all the data that a sensor will produce at every discretized time interval. Reading a value from a sensor is analogous to extracting a value from this array indexed at the current time (t). We define an abstract matrix \mathcal{E} to correspond to the physical sensing environment such that $\mathcal{E}_{i,j}$ is the reading (value) from sensor i at discretized time j . We define the function $f_{\mathcal{E}}(i, j)$ to extract the j^{th} value (corresponding with time j) from stream (sensor) i from \mathcal{E} .

We define the (strict) function `get` to extract a value from the sensing environment. In the simplified Core STEP we support only two types of sensors, a time sensor (`time`) and Image sensors (`image`).

$$(E-GET1) \frac{e_1 \mid \nu \mid t \rightarrow e'_1 \mid \nu' \mid t + 1}{\mathit{get} \ e_1 \mid \nu \mid t \rightarrow \mathit{get} \ e'_1 \mid \nu' \mid t + 1}$$

$$(E-GETAPPLY) \quad \mathit{get} \ v_1 \mid \nu \mid t \rightarrow v_2 \mid \nu \mid t + 1$$

where $v_2 = f_{\mathcal{E}}(v, t)$
and $f_{\mathcal{E}} : \mathit{Sensor} \ \tau \times \mathit{Int} \rightarrow \tau$

The `time` sensor returns the number of evaluations (computations) since the beginning of the evaluation (*i.e.*, returns t).

$$(E-GETTIME) \quad (\mathit{get} \ \mathit{time} \mid \nu \mid t) \rightarrow (t \mid \nu \mid t + 1)$$

3.3.4 Let binding (by-need with caching)

Again, our let semantic is a hybrid approach that is deferred evaluation (*ala lazy*) coupled with evaluation re-use (*ala by-value/eager*). To facilitate this, when a let is encountered, its assignments are added directly to the store ν (*via* E-LETN) without evaluation. Expressions and instead are evaluated within ν when their variable is encountered elsewhere (*via* E-VAR1 and E-VAR2) including variables that point to other variables. Our let expression allows for simultaneous assignment; in E-LETN we assume that all terms x_i are assigned simultaneously and may have interdependency.³

Variable terms are subject to alpha renaming to avoid variable capture, *etc.*

$$(E\text{-LETN}) \quad \frac{\text{let } \{x_1 = \mathbf{e}_1, \dots, x_n = \mathbf{e}_n\} \text{ in } \mathbf{e}_{n+1} \mid \nu \mid t}{\rightarrow \mathbf{e}_{n+1} \mid \nu \cup \{x_1 \mapsto \mathbf{e}_1, \dots, x_n \mapsto \mathbf{e}_n\} \mid t + 1}$$

$$(E\text{-VAR1}) \quad \frac{\mathbf{e}_1 \mid \nu \mid t \rightarrow \mathbf{e}'_1 \mid \nu' \mid t + 1}{(x \mid \nu \cup \{x \mapsto \mathbf{e}_1\} \mid t) \rightarrow (x \mid \nu' \cup \{x \mapsto \mathbf{e}'_1\} \mid t + 1)} \quad x \notin \text{domain}(\nu) \cup \text{domain}(\nu')$$

$$(E\text{-VAR2}) \quad (x \mid \nu \cup \{x \mapsto \mathbf{v}\} \mid t) \rightarrow (\mathbf{v} \mid \nu \cup \{x \mapsto \mathbf{v}\} \mid t + 1)$$

3.3.5 Triggers

STEP provides a **trigger** construct to specify repetitive conditional evaluation. It repeatedly evaluates a boolean expression until it is true and then evaluates a second expression (the first expression “triggers” the second). We also provide the ability to provide sequential let-bindings that behave as though they are within the scope of each trigger expression evaluation (*i.e.*, the let-term is recomputed on every expansion of the trigger) are available to both branches of the trigger (a single let-binding result spans both trigger arguments).

We define the trigger’s repetition recursively, via the conditional. For completeness we also define functionally degenerate trigger forms; expressions that specify a constant value for the trigger predicate produce a superfluous trigger expression (*i.e.*, **trigger true** \mathbf{e}_1 always reduces to \mathbf{e}_1 while **trigger false** \mathbf{e}_1 would never proceed).

$$(E\text{-TRIGGER1}) \quad \frac{\text{trigger } \mathbf{e}_1 \ \mathbf{e}_2 \mid \nu \mid t}{\rightarrow \text{cond } \mathbf{e}_1 \ \mathbf{e}_2 \ (\text{trigger } \mathbf{e}_1 \ \mathbf{e}_2) \mid \nu \mid t + 1}$$

$$(E\text{-TRIGGERDEG}) \quad \frac{\mathbf{e}_2 \mid \nu \mid t \rightarrow \mathbf{e}'_2 \mid \nu' \mid t + 1}{\text{trigger } \mathbf{v} \ \mathbf{e}_2 \mid \nu \mid t \rightarrow \text{cond } \mathbf{v} \ \mathbf{e}'_2 \ (\text{trigger } \mathbf{v} \ \mathbf{e}_2) \mid \nu' \mid t + 1}$$

The **letonce** expression shown is syntactic sugar for the expanded trigger expression, which builds a sequential set of variable assignments for which there is one one expansion per trigger iteration (alpha renaming ensures we are not using the same variable in every iteration). The binding is intentionally placed at the scope of both the predicate and the conclusion (*i.e.*, consequent) of the trigger term.

$$\text{letonce } \{x_1 = \mathbf{e}_1, \dots, x_n = \mathbf{e}_n\} \text{ in trigger } \mathbf{e}_{n+1} \ \mathbf{e}_{n+2} \equiv \text{trigger } \{x_1 = \mathbf{e}_1, \dots, x_n = \mathbf{e}_n\} \ \mathbf{e}_{n+1} \ \mathbf{e}_{n+2}$$

³Put differently, the multiple assignments in E-LETN are *not* syntactic sugar for nested lets.

$$\begin{array}{l}
\text{trigger } \{x_1=e_1, \dots, x_n=e_n\} \ e_{n+1} \ e_{n+2} \\
\text{(E-TRIGGERLET)} \quad \rightarrow \text{let } \{x_1=e_1, \dots, x_n=e_n\} \ \text{in} \\
\quad \text{cond } e_{n+1} \ e_{n+2} \ (\text{trigger } \{x_1=e_1, \dots, x_n=e_n\} \ e_{n+1} \ e_{n+2})
\end{array}$$

3.4 Syntax of Types

The syntax of types in Core STEP are given below.

$$\begin{array}{ll}
t ::= & \text{base types} \\
& Int \mid Bool \mid Img \\
\tau_p ::= & \text{primitive types} \\
& t^{\{s,s\}} \quad w/ \text{ size annotation} \\
\tau ::= & \text{types} \\
& \tau_p \quad \text{primitive type} \\
& Sensor \ \tau_b \quad \text{sensor} \\
& \tau \rightarrow \tau \quad \text{opcode} \\
s ::= & n \mid r \quad \text{size annotation}
\end{array}$$

3.5 Static Semantics and Sizing of Types

The typing (static semantic) of an expression takes the general form:

$$\Gamma \vdash e : t^{\{s_{min}, s_{max}\}} \ \$c, \kappa$$

Where t is base type (*e.g.*, Int , $Bool$, Img), $\{s_{min}, s_{max}\}$ is the *size* annotation for the type (s_{min} is the lower size bound, s_{max} is the upper size bound)⁴, c is the worst-case approximation of computational cost of the expression and κ is a size constraint set (s_{max} and s_{min} are size annotations constrained by the simple equations stored in κ , as we will see later). In English we read this as: “Expression e and has a worst-case computational cost of c and is of type t under the typing environment Γ , where t has a minimum size s_{min} , a maximum size s_{max} , and is subject to size constraints κ .”

For clarity of presentation we denote size annotations with different symbols depending on the base type they annotate. We use n_i to represent a size annotation for an Integer, and r_i for a size annotation on an Image. In our presentation we omit the implicit constraint $s_1 \leq s_2$ present for every size vector $\{s_1, s_2\}$,

3.5.1 Primitive Types

The typing rules for values are given below. The computational cost (c) for a value is always 0.

$$\text{(T-INT)} \quad \frac{}{\Gamma \vdash \mathbf{n} : Int^{\{n,n\}} \ \$0, \{n = \mathbf{n}\}}$$

A constant integer value has both size annotation variables constrained to the integer’s actual value.

⁴This pair can be expanded to a tuple to track more dimensions/aspects, as described in Section 5

$$(T\text{-IMAGE}) \frac{}{\Gamma \vdash \mathbf{i} : \text{Img}^{\{r,r\}} \ \$ \ 0, \{r = \text{resolutionof}(\mathbf{i})\}}$$

The size of an image is given by its resolution, taken logically to be the number of pixels in the Image, however to simplify our presentation we use only the width of the image. The pair specifies the range of possible resolutions for the image; As with integers, both values will be the same for a specific, concrete instantiation of an Image.

$$(T\text{-BOOL}) \frac{}{\Gamma \vdash \mathbf{b} : \text{Bool} \ \$ \ 0, \emptyset}$$

A boolean (`true` | `false`) has a negligible fixed size and thus its size annotation is omitted (*i.e.*, $\{1,1\}$).

We define some convenience functions to manipulate the type and its size annotation :

$$\begin{aligned} \text{minsize}(t^{\{n_1, n_2\}}) &= n_1 & \text{maxsize}(t^{\{n_1, n_2\}}) &= n_2 \\ \text{base}(t^{\{n_1, n_2\}}) &= t \end{aligned}$$

For example: $\text{minsize}(\text{Int}^{\{4,8\}}) = 4$ $\text{base}(\text{Int}^{\{4,8\}}) = \text{Int}$ $\text{maxsize}(\text{Img}^{\{2,n\}}) = n$

3.5.2 Subtyping (bounding sizes)

$$\begin{aligned} (S\text{-REFL}) \quad \tau <: \tau & & (S\text{-TRANS}) \quad \frac{\tau_1 <: \tau_2 \quad \tau_2 <: \tau_3}{\tau_1 <: \tau_3} \\ (S\text{-ARROW}) \quad \frac{\tau_2 <: \tau_1 \quad \tau'_1 <: \tau'_2}{\tau_1 \rightarrow \tau'_1 <: \tau_2 \rightarrow \tau'_2} & & (S\text{-PAIR}) \quad \frac{\tau_1 <: \tau_2 \quad \tau'_1 <: \tau'_2}{\{\tau_1 \times \tau'_1\} <: \{\tau_2 \times \tau'_2\}} \\ (S\text{-SENSOR}) \quad \frac{\tau_1 <: \tau_2}{\text{Sensor } \tau_1 <: \text{Sensor } \tau_2} \end{aligned}$$

We define a relation similar to [LH96]’s subtyping relation (\sqsubseteq) which allows a weakening of the type to increase a size bound, in order to provide an upper bound of the size of the input relative to work to be completed. In our environment we notice that the correctness of Image processing functions may be impacted by the size of the input (*i.e.*, resolution of the image); as such we cannot arbitrarily increase the logical size (resolution) of this data without adverse consequences to functional correctness.

The need to track the lower bound extends into all aspects of the sized typing system, so we use a general sizing/weaken rule (S-SIZED) to describe the subtype relationship for specific sized types.

$$(S\text{-SIZED}) \quad \frac{(s_{min} \geq s'_{min}) \quad (s_{max} \leq s'_{max})}{t^{\{s_{min}, s_{max}\}} <: t^{\{s'_{min}, s'_{max}\}}}$$

Finally, we give the rule for weakening via the subtype relation, which should be clear when considered with S-SIZED, above. Notice the constraint set κ grows to include the sizing relationship between τ_1 and τ_2 . If one were to expand the sizing pair to include more dimensions/aspects of type size, then the S-SIZED and T-WEAKEN constraints would be augmented to support the new sizing logic (Section 5 has more on this).

$$(T\text{-WEAKEN}) \quad \frac{\Gamma \vdash \mathbf{e} : \tau_1 \ \$ \ c, \kappa \quad \tau_1 <: \tau_2}{\Gamma \vdash \mathbf{e} : \tau_2 \ \$ \ c, \kappa \cup \kappa_2}$$

where $\kappa_2 = \{\text{minsize}(\tau_2) \leq \text{minsize}(\tau_1), \text{maxsize}(\tau_2) \geq \text{maxsize}(\tau_1)\}$

3.5.3 Conditional

$$(T\text{-COND}) \quad \frac{\Gamma \vdash \mathbf{e}_1 : Bool \ \$ \ c_1, \kappa_1 \quad \Gamma \vdash \mathbf{e}_2 : \tau \ \$ \ c_2, \kappa_2 \quad \Gamma \vdash \mathbf{e}_3 : \tau \ \$ \ c_3, \kappa_3}{\Gamma \vdash \mathbf{cond} \ \mathbf{e}_1 \ \mathbf{e}_2 \ \mathbf{e}_3 : \tau \ \$ \ 1 + c_1 + \max(c_2 \ c_3), \kappa_1 \cup \kappa_2 \cup \kappa_3}$$

In the (common) event where terms of the conditional branches are different sizes, the application of weaken can be used to relax the bounds on either side to meet at the lower minimum size and larger maximum size.

In the event that either branch's type contains a size variable, each branch may be weakened to a new, common size variable for the conditional. The example that follows portrays exactly this, in which two images (or expressions of type image) that have different, yet unknown sizes that are supplied as the branches of a conditional only after applying T-WEAKEN to each to arrive at the new size variables n_5 and n_6 :

$$\frac{\begin{array}{c} \vdots \\ \mathbf{b} : Bool \ \$c_0, \kappa_0 \end{array} \quad \frac{\mathbf{i}_1 : Img \ \{r_1, r_2\} \ \$ \ c_1, \kappa_1}{\mathbf{i}_1 : Img \ \{r_5, r_6\} \ \$ \ c_1, \kappa_1 \cup \kappa'_1} (T\text{-WEAKEN}) \quad \frac{\mathbf{i}_2 : Img \ \{r_3, r_4\} \ \$ \ c_2, \kappa}{\mathbf{i}_2 : Img \ \{r_5, r_6\} \ \$ \ c_2, \kappa \cup \kappa'_2} (T\text{-WEAKEN})}{\mathbf{cond} \ \mathbf{b} \ \mathbf{i}_1 \ \mathbf{i}_2 : Img \ \{r_5, r_6\} \ \$ \ c_0 + \max(c_1, c_2), \kappa_0 \cup \kappa_1 \cup \kappa'_1 \cup \kappa_2 \cup \kappa'_2} (T\text{-COND})$$

where:

$$\begin{aligned} \kappa'_1 &= \{r_5 \leq r_1, r_6 \geq r_2\} \\ \kappa'_2 &= \{r_5 \leq r_3, r_6 \geq r_4\} \end{aligned}$$

3.5.4 Sensors

A Sensor is a “container” type; a sensor of type *Sensor* τ will return values of type τ when “read”. The Sensor type has an negligible, omitted static size, however the inner (contained) type can be annotated with size bounds to indicate the capabilities of the sensor (*e.g.*, the range of resolutions a camera can support). The rule below indicates that `image` is an image sensor that can return images ranging in size from \mathbf{r}_{min} to \mathbf{r}_{max} (that should correlate with the maximum and minimum resolution capabilities of physical hardware). As this judgment has no premise and introduces the size variables r_1 and r_2 into the derivation tree, we will solve our constraint set for these variables once type derivation is complete.

$$(T\text{-IMAGESENSOR}) \quad \frac{}{\Gamma \vdash \mathbf{image} : Sensor \ Img \ \{r_1, r_2\} \ \$ \ 0, \{r_1 \geq \mathbf{r}_{min}\} \cup \{r_2 \leq \mathbf{r}_{max}\}}$$

The `time` sensor returns the number of evaluations (computations) since the beginning of the evaluation.

$$(T\text{-TIMESENSOR}) \quad \frac{}{\Gamma \vdash \mathbf{time} : Sensor \ Int \ \{n_1, n_2\} \ \$ \ 0, \{n_1 \geq 0\}}$$

$$(T\text{-SENSORREAD}) \quad \frac{\mathbf{type}(\mathbf{get}) = Sensor \ \tau_1 \rightarrow \tau_1 \quad \Gamma \vdash \mathbf{e}_1 : Sensor \ \tau_1 \ \$ \ c, \kappa_1}{\Gamma \vdash \mathbf{get} \ \mathbf{e}_1 : \tau_1 \ \$ \ c_1 + 1 + \mathbf{latentcost}(\mathbf{get}, \mathbf{maxsize}(\tau_1)), \kappa}$$

3.5.5 Opcodes

$$(T\text{-OP}) \quad \frac{\mathbf{type}(op) = \tau_1 \rightarrow \tau_2 \quad \mathbf{constr}(op) = \kappa}{\Gamma \vdash op : \tau_1 \rightarrow \tau_2 \ \$ \ 0, \kappa}$$

$$(T\text{-OPAPPLY}) \frac{\Gamma \vdash op : \tau_1 \rightarrow \tau_2 \ \$ \ 0, \kappa_1 \quad \Gamma \vdash e_2 : \tau_1 \ \$ \ c, \kappa_2}{\Gamma \vdash op \ e_2 : \tau_2 \ \$ \ 1 + c + \text{latentcost}(op, \text{maxsize}(\tau_1)), \kappa_1 \cup \kappa_2}$$

where $op \in \{op_1 \mid op_2 \mid op_3 \mid \dots\}$

The $\text{latentcost}()$ function [RG94] returns the discretized computational cost of each opcode as an equation of the size of its input. For example, the complexity of finding a face in an image is a function of the total number of pixels in the image.

3.5.6 Let binding

An instance of a variable has whatever type is assumed for it in the typing environment Γ . A variable has no computational cost or constraints associated with it, rather the costs and constraints are assigned when variables are bound (*i.e.*, in the **let** term).

$$(T\text{-VAR}) \frac{\Gamma(x) = \tau}{\Gamma \vdash x : \tau \ \$ \ 0, \emptyset}$$

In a **let** term, we take the sum of the costs of the let-bound expressions as well as the union of all of their associated constraints.

$$(T\text{-LETN}) \frac{\Gamma \vdash e_1 : \tau_1 \ \$ \ c_1, \kappa_1 \ \dots \ \Gamma \vdash e_n : \tau_n \ \$ \ c_n, \kappa_n \quad \Gamma, x_1 : \tau_1, \dots, x_n : \tau_n \vdash e_{n+1} : \tau_{n+1} \ \$ \ c_{n+1}, \kappa_{n+1}}{\Gamma \vdash \text{let } \{x_1 = e_1, \dots, x_n = e_n\} \ \text{in } e_{n+1} : \tau_{n+1} \ \$ \ c_{n+1} + \sum_{i=1}^n c_i, \kappa_{n+1} \cup \bigcup_{i=1}^n \kappa_i}$$

The cost defined in this rule is an over-estimate of the total cost as the let-bound symbols (some x_i) may not occur in the evaluation path of e_{n+1} . These costs will not, however, be “charged” twice, as a variable itself has cost zero when computing the cost of e_{n+1} .

3.5.7 Triggers

A trigger’s cost and constraint takes much the same form as the **cond** and **let** instances from which the trigger is derived.

$$(T\text{-TRIG}) \frac{\Gamma \vdash e_i : \tau_i \ \$ \ c_i, \kappa_i \ (\text{for } i \in 1..n) \quad \Gamma, \{x_i : \tau_i\}^{1..n} \vdash e_{n+1} : \text{Bool} \ \$ \ c_{n+1}, \kappa_{n+1} \quad \Gamma, \{x_i : \tau_i\}^{1..n} \vdash e_{n+2} : \tau \ \$ \ c_{n+2}, \kappa_{n+2}}{\Gamma \vdash \text{trigger } \{x_i = e_i\}^{1..n} e_{n+1} \ e_{n+2} : \tau \ \$ \ \mathcal{T} * (c_{n+1} + \sum_{i=1}^n c_i) + c_{n+2}, \kappa_{n+1} \cup \kappa_{n+2} \cup \bigcup_{i=1}^n \kappa_i}$$

where \mathcal{T} is a new cost variable from the set of unused cost variables, \mathcal{C}

Specific values for \mathcal{T} may be provided by explicit user bounds or other means.

3.5.8 Typing of the Let-store, ν

We can assign a type (more accurately sequence of types) to a variable store ν if we have assumed types for all of the variables contained within the store. Similarly we associate a cost and constraint set with the store (for use in interim steps of typing derivations).

$$(T\text{-NU}) \frac{\Gamma, \Gamma' \vdash \nu(x_i) : \Gamma'(x_i) \ \$ \ c_i, \kappa_i \quad (\text{for every } x_i \in \{x_1, \dots, x_n\} = \text{Domain}(\nu))}{\Gamma \vdash \nu : \Gamma' \ \$ \ \sum_{i=1}^n c_i, \bigcup_{i=1}^n \kappa_i}$$

Finally we can combine an expression *and* a typed variable store into a typed pair by discharging the typing assumptions of the variable store, as shown below.

$$(T\text{-COMPLETE}) \quad \frac{\Gamma \vdash \mathbf{e} : \tau \ \$ \ c_1, \kappa_1 \quad \emptyset \vdash \nu : \Gamma \ \$ \ c_2, \kappa_2}{\emptyset \vdash \mathbf{e} \mid \nu : \tau \ \$ \ c_1 + c_2, \kappa_1 \cup \kappa_2}$$

3.6 Soundness of Core STEP

This Section proves *Soundness* for the Core STEP (*Soundness = Progress + Preservation*). *Progress* means that every expression is either a value or can take an evaluative step (*i.e.*, expressions don't get stuck), while *Preservation* means that every typed expression that takes an evaluative step results in another typed expression (*i.e.*, evaluation is type preserving). We begin by proving some Lemmas that will be useful for our proofs of Progress and Preservation.

3.6.1 Lemma: L-Var

Suppose $\emptyset \vdash \mathbf{e} \mid \nu : \tau \ \$ \ c, \kappa$ and $\mathbf{e} \mid \nu$ is closed.

If $\mathbf{e} = x_i$, then for some c' and κ' :

$$\emptyset \vdash \nu(x_i) \mid \nu : \tau \ \$ \ c', \kappa'$$

Proof: By the structure of the derivation of $\emptyset \vdash \mathbf{e} \mid \nu : \tau \ \$ \ c, \kappa$ where $\mathbf{e} = x_i$.

$$\frac{\frac{\Gamma(x_i) = \tau}{\Gamma \vdash x_i : \tau \ \$ \ 0, \emptyset} (T\text{-VAR}) \quad \frac{\Gamma \vdash \nu(x_j) : \Gamma(x_j) \ \$ \ c_j, \kappa_j \quad (\text{for every } x_j \in \{x_1, \dots, x_n\} = \text{Domain}(\nu))}{\emptyset \vdash \nu : \Gamma \ \$ \ c, \kappa} (T\text{-NU})}{\emptyset \vdash x_i \mid \nu : \Gamma(x_i) \ \$ \ c, \kappa} (T\text{-COMPLETE})$$

From the premise of T-NU: $\Gamma \vdash \nu(x_i) : \Gamma(x_i) \ \$ \ c_i, \kappa_i$

From the premise of T-VAR: $\Gamma(x_i) = \tau$

Combining these we have: $\Gamma \vdash \nu(x_i) : \tau \ \$ \ c_i, \kappa_i$

$$\frac{\Gamma \vdash \nu(x_i) : \tau \ \$ \ c_i, \kappa_i \quad \emptyset \vdash \nu : \Gamma \ \$ \ c, \kappa}{\emptyset \vdash \nu(x_i) \mid \nu : \tau \ \$ \ c', \kappa'} (T\text{-COMPLETE})$$

3.6.2 Lemma(s): Inversion+ ν

Suppose $\emptyset \vdash \mathbf{e} \mid \nu : \tau \ \$ \ c, \kappa$ is the last judgment \mathcal{J}_0 in a typing derivation tree \mathcal{D} . Then:

- (1) The left premise of this judgment in \mathcal{D} is $\Gamma \vdash \mathbf{e} : \tau \ \$ \ c'', \kappa''$ (say \mathcal{J}_1) for some c'' and κ''
- (2) If \mathcal{J}_2 is a premise of \mathcal{J}_1 of the form $\Gamma' \vdash \mathbf{e}' : \tau' \ \$ \ c', \kappa'$ for some c', κ' then $\Gamma' \vdash \mathbf{e}' \mid \nu : \tau' \ \$ \ c', \kappa'$ can be derived for all such \mathcal{J}_2 .

Proof: (Generic) By the structure of the derivation of $\emptyset \vdash \mathbf{e} \mid \nu : \tau \ \$ \ c, \kappa$ (below). In the general form, to have reached $\emptyset \vdash \mathbf{e} \mid \nu : \tau \ \$ \ c, \kappa$, we must have an application of T-COMPLETE with right-side premise $\emptyset \vdash \nu : \Gamma \ \$ \ c'', \kappa''$ and on the left side of the derivation we find the individual sub-premisses to type the expression \mathbf{e} . We call this rule *T-Rule* as a placeholder for specific instances of \mathbf{e} and similarly denote the sub-premisses as \mathcal{J}_2 and the consequent (in which \mathbf{e} is given a type) as \mathcal{J}_1 . We can apply

T-COMPLETE with the right hand side premise of the existing T-COMPLETE to each sub-premise of \mathcal{J}_2 individually to arrive at our conclusion. We detail the specific individual cases below.

$$\frac{\mathcal{J}_2}{\mathcal{J}_1 = \Gamma \vdash \mathbf{e} : \tau \$ c''', \kappa'''} \text{(T-Rule)} \frac{\Gamma \vdash \nu(x_j) : \Gamma(x_j) \$ c_j, \kappa_j \text{ (for every } x_j \in \{x_1, \dots, x_n\} = \text{Domain}(\nu))}{\emptyset \vdash \nu : \Gamma \$ c'', \kappa''} \text{(T-NU)}}{\mathcal{J}_0 = \emptyset \vdash \mathbf{e} \mid \nu : \tau \$ c, \kappa} \text{(T-COMPLETE)}$$

Case L-COND:

Suppose $\Gamma \vdash \mathbf{e} \mid \nu : \tau \$ c, \kappa$ and $\mathbf{e} \mid \nu$ is closed.

If $\mathbf{e} = \text{cond } \mathbf{e}_1 \ \mathbf{e}_2 \ \mathbf{e}_3$ then for some c'_1, c'_2, c'_3 and $\kappa'_1, \kappa'_2, \kappa'_3$:

$$\emptyset \vdash \mathbf{e}_1 \mid \nu : \text{Bool} \$ c'_1, \kappa'_1$$

$$\emptyset \vdash \mathbf{e}_2 \mid \nu : \tau \$ c'_2, \kappa'_2$$

$$\emptyset \vdash \mathbf{e}_3 \mid \nu : \tau \$ c'_3, \kappa'_3$$

Proof: Using the generic proof above, where:

$$\mathcal{J}_2 = \Gamma \vdash \mathbf{e}_1 : \text{Bool} \$ c_1, \kappa_1 \quad \Gamma \vdash \mathbf{e}_2 : \tau \$ c_2, \kappa_2 \quad \Gamma \vdash \mathbf{e}_3 : \tau \$ c_3, \kappa_3$$

$$\mathcal{J}_1 = \Gamma \vdash \text{cond } \mathbf{e}_1 \ \mathbf{e}_2 \ \mathbf{e}_3 : \tau \$ (1 + c_1 + (\max)(c_2, c_3)), (\kappa_1 \cup \kappa_2 \cup \kappa_3)$$

T-Rule = T-COND.

Case L-GET:

Suppose $\Gamma \vdash \mathbf{e} \mid \nu : \tau \$ c, \kappa$ and $\mathbf{e} \mid \nu$ is closed.

If $\mathbf{e} = \text{get } \mathbf{e}_1$ then for some c'_1 and κ'_1 :

$$\emptyset \vdash \mathbf{e}_1 \mid \nu : \text{Sensor } \tau_1 \$ c'_1, \kappa'_1$$

Proof: Using the generic proof above, where:

$$\mathcal{J}_2 = \Gamma \vdash \mathbf{e}_1 : \text{Sensor } \tau_1 \$ c_1, \kappa_1$$

$$\mathcal{J}_1 = \Gamma \vdash \text{get } \mathbf{e}_1 : \tau_1 \$ c'_1, \kappa'_1$$

T-Rule = T-SENSORREAD.

Case L-OPAPPLY:

Suppose $\Gamma \vdash \mathbf{e} \mid \nu : \tau \$ c, \kappa$ and $\mathbf{e} \mid \nu$ is closed.

If $\mathbf{e} = \text{op } \mathbf{e}_2$ then for some c'_1, c'_2 and κ'_1, κ'_2 :

$$\emptyset \vdash \text{op} \mid \nu : \tau_1 \rightarrow \tau_2 \$ c'_1, \kappa'_1$$

$$\emptyset \vdash \mathbf{e}_2 \mid \nu : \tau_1 \$ c'_2, \kappa'_2$$

Proof: Using the generic proof above, where:

$$\mathcal{J}_2 = \Gamma \vdash \text{op} : \tau_1 \rightarrow \tau_2 \$ 0, \kappa_1 \quad \Gamma \vdash \mathbf{e}_2 : \tau_1 \$ c, \kappa_2$$

$$\mathcal{J}_1 = \Gamma \vdash \text{op } \mathbf{e}_2 : \tau \$ c, \kappa_1 \cup \kappa_2$$

T-Rule = T-OPAPPLY.

Case L-LETN:

Suppose $\Gamma \vdash \mathbf{e} \mid \nu : \tau \$ c, \kappa$ and $\mathbf{e} \mid \nu$ is closed.

If $\mathbf{e} = \text{let } \{x_1 = \mathbf{e}_1, \dots, x_n = \mathbf{e}_n\} \text{ in } \mathbf{e}_{n+1}$ then for some $c'_1 \dots c'_{n+1}$ and $\kappa'_1 \dots \kappa'_{n+1}$:

$$\emptyset \vdash \mathbf{e}_1 \mid \nu : \tau_1 \$ c'_1, \kappa'_1$$

\vdots

$$\emptyset \vdash \mathbf{e}_n \mid \nu : \tau_n \$ c'_n, \kappa'_n$$

$$\emptyset, x_1 : \tau_1, \dots, x_n : \tau_n \vdash \mathbf{e}_{n+1} \mid \nu : \tau_{n+1} \$ c'_{n+1}, \kappa'_{n+1}$$

Proof: Using the generic proof above, where:

$$\begin{aligned} \mathcal{J}_2 &= \Gamma \vdash \mathbf{e}_1 : \tau_1 \ \$ \ c_1, \kappa_1 \ \dots \ \Gamma \vdash \mathbf{e}_n : \tau_n \ \$ \ c_n, \kappa_n \ \ \Gamma, x_1 : \tau_1, \dots, x_n : \tau_n \vdash \mathbf{e}_{n+1} : \\ &\tau_{n+1} \ \$ \ c_{n+1}, \kappa_{n+1} \\ \mathcal{J}_1 &= \Gamma \vdash \mathbf{let} \ \{x_1 = \mathbf{e}_1, \dots, x_n = \mathbf{e}_n\} \ \mathbf{in} \ \mathbf{e}_{n+1} : \tau \ \$ \ c_{n+1} + \sum_{i=1}^n c_i, \kappa_{n+1} \cup \bigcup_{i=1}^n \kappa_i \\ \text{T-Rule} &= \text{T-LETN}. \end{aligned}$$

Case L-TRIGGER:

Suppose $\Gamma \vdash \mathbf{e} \mid \nu : \tau \ \$ \ c, \kappa$ and $\mathbf{e} \mid \nu$ is closed.

If $\mathbf{e} = \mathbf{trigger} \ \{x_i = \mathbf{e}_i\}^{i \in 1..n} \mathbf{e}_{n+1} \ \mathbf{e}_{n+2}$ then for some $c'_1 \dots c'_{n+2}$ and $\kappa'_1 \dots \kappa'_{n+2}$:

$$\begin{aligned} \emptyset \vdash \mathbf{e}_1 \mid \nu : \tau_1 \ \$ \ c'_1, \kappa'_1 \\ \dots \\ \emptyset \vdash \mathbf{e}_n \mid \nu : \tau_n \ \$ \ c'_n, \kappa'_n \\ \emptyset, \{x_i : \tau_i\}^{1..n} \vdash \mathbf{e}_{n+1} \mid \nu : \text{Bool} \ \$ \ c'_{n+1}, \kappa'_{n+1} \\ \emptyset, \{x_i : \tau_i\}^{1..n} \vdash \mathbf{e}_{n+2} \mid \nu : \tau \ \$ \ c'_{n+2}, \kappa'_{n+2} \end{aligned}$$

Proof: Using the generic proof above, where:

$$\begin{aligned} \mathcal{J}_2 &= \Gamma \vdash \mathbf{e}_1 : \tau_1 \ \$ \ c_1, \kappa_1 \ \dots \ \Gamma \vdash \mathbf{e}_n : \tau_n \ \$ \ c_n, \kappa_n \ \ \Gamma, \{x_i : \tau_i\}^{1..n} \vdash \mathbf{e}_{n+1} : \\ &\text{Bool} \ \$ \ c_{n+1}, \kappa_{n+1} \ \ \Gamma, \{x_i : \tau_i\}^{1..n} \vdash \mathbf{e}_{n+2} : \tau \ \$ \ c_{n+2}, \kappa_{n+2} \\ \mathcal{J}_1 &= \Gamma \vdash \mathbf{trigger} \ \{x_i = \mathbf{e}_i\}^{i \in 1..n} \mathbf{e}_{n+1} \ \mathbf{e}_{n+2} : \tau \ \$ \ \mathcal{T} * (c_{n+1} + \sum_{i=1}^n c_i) + c_{n+2}, \kappa_{n+1} \cup \\ &\kappa_{n+2} \cup \bigcup_{i=1}^n \kappa_i \\ \text{T-Rule} &= \text{T-TRIG}. \end{aligned}$$

3.6.3 Theorem [Progress]

Suppose $\mathbf{e} \mid \nu$ is closed and $\emptyset \vdash \mathbf{e} \mid \nu : \tau \ \$ \ c, \kappa$ (for some cost c and constraint set κ)

Either: (1) \mathbf{e} is a value or (2) there exists some \mathbf{e}' and store ν' such that for every

$$\begin{aligned} t: \\ (\mathbf{e} \mid \nu \mid t) \rightarrow (\mathbf{e}' \mid \nu' \mid t+1). \end{aligned}$$

Proof: By induction on the number of *unique* sub-derivations of a typing derivation of $\emptyset \vdash \mathbf{e} \mid \nu : \tau \ \$ \ c, \kappa$.

We proceed by analysis of the shape of \mathbf{e} to show that the property holds for the larger derivation.

Case $\mathbf{e} = \mathbf{v}$:

Satisfied trivially as these terms are values.

Case $\mathbf{e} = \mathbf{cond} \ \mathbf{e}_1 \ \mathbf{e}_2 \ \mathbf{e}_3$:

Applying the Induction Hypothesis to a derivation whose final judgment is $\Gamma \vdash \mathbf{e}_1 \mid \nu : \text{Bool} \ \$ \ c_1, \kappa_1$ (via L-COND) we can verify that progress holds for $\mathbf{e}_1 \mid \nu$ (*i.e.*, \mathbf{e}_1 is a value or $\mathbf{e}_1 \mid \nu$ can take an evaluative step). So either (1) \mathbf{e}_1 is a value (specifically, a Boolean) and E-IFTRUE or E-IFFALSE can be applied or $\mathbf{e}_1 \mid \nu$ can take an evaluative step such that E-IF applies and $(\mathbf{e} \mid \nu \mid t) \rightarrow (\mathbf{cond} \ \mathbf{e}'_1 \ \mathbf{e}_2 \ \mathbf{e}_3 \mid \nu' \mid t+1)$.

Case $\mathbf{e} = x_i$ with $x_i : \tau_i$ and $x_i \in \text{Domain}(\nu)$ (as $\mathbf{e} \mid \nu$ is closed):

By the definition of ν (and closure of the pair $\mathbf{e} \mid \nu$) we re-write $\emptyset \vdash x_i \mid \nu : \tau$ as $\emptyset \vdash x_i \mid \{x_i \mapsto \mathbf{e}_i\} \cup \nu_+ : \tau$. By the definition of ν either $\mathbf{e}_i = \mathbf{v}_i$ and E-VAR applies without premise or \mathbf{e}_i is a composite expression. To apply E-VAR1 we require progress for $\mathbf{e}_i \mid \nu$. By the

definition of ν we can re-write e_i as $\nu(x_i)$ for which progress holds via L-VAR (to which we apply the the Inductive Hypothesis). Hence E-VAR or E-VAR1 may be applied and progress holds for this term.

Case $e = \text{let } \{x_1 = e_1, \dots, x_n = e_n\} \text{ in } e_{n+1}$:

E-LETN takes an evaluative step without premise such that $(e \mid \nu \mid t) \rightarrow (e_{n+1} \mid \nu \cup \{x_1 \mapsto e_1, \dots, x_n \mapsto e_n\} \mid t + 1)$.

Case $e = \text{trigger } \{x_1=e_1, \dots, x_n=e_n\} e_{n+1} e_{n+2}$:

E-TRIGGERLET applies without premise for a single step of evaluation.

In the event that the let assignment set is empty (*i.e.*, $e = \text{trigger } e_{n+1} e_{n+2}$) then E-TRIGGER1 applies without premise.

Case $e = \text{get } e_1$:

Either e_1 is a value and E-GETAPPLY applies (or E-GETTIME if $e_1 = \text{time}$), or by L-GET and the Inductive Hypothesis, we show progress holds on $(e_1 \mid \nu)$ so $e_1 \mid \nu \mid t \rightarrow e'_1 \mid \nu' \mid t + 1$ and E-GET1 applies such that $(e \mid \nu \mid t) \rightarrow (\text{get } e'_1 \mid \nu' \mid t + 1)$

Case $e = \text{op } e_2$:

Either e_2 is a value and E-OPAPPLY applies, or by L-OPAPPLY and the Inductive Hypothesis, progress holds on $e_2 \mid \nu$ so $e_2 \mid \nu \mid t \rightarrow e'_2 \mid \nu' \mid t + 1$ and E-OP1 applies such that $(e \mid \nu \mid t) \rightarrow (\text{op } e'_2 \mid \nu' \mid t + 1)$

3.6.4 Theorem [Preservation]

Suppose $e \mid \nu$ is closed and $\emptyset \vdash e \mid \nu : \tau \$ c, \kappa$.

If $(e \mid \nu \mid t) \rightarrow (e' \mid \nu' \mid t + 1)$, then for some cost c' and constraint set κ' :

$$\emptyset \vdash e' \mid \nu' : \tau \$ c', \kappa'$$

Proof: By induction on the number of *distinct* sub-derivations of a typing derivation of $\emptyset \vdash e \mid \nu : \tau \$ c, \kappa$.

We proceed by analysis of the shape of e to show that the property holds for the larger derivation.

Case $e = v$:

These cases are satisfied trivially, as these terms are values and do not make a further evaluative step.

Case $e = \text{cond } e_1 e_2 e_3$:

We consider each possible evaluative case, individually.

Case E-IFTRUE:

$$(\text{cond true } e_2 e_3 \mid \nu \mid t) \rightarrow (e_2 \mid \nu \mid t + 1)$$

We consider the typing derivation of $e \mid \nu$ in which we have $\Gamma \vdash e_2 : \tau \$ c_2, \kappa_2$ (from the premise T-COND) and $\nu : \Gamma \$ c'', \kappa''$ (from the premise T-COMPLETE). We combine them in a new application of T-COMPLETE to verify $\emptyset \vdash e_2 \mid \nu : \tau \$ c', \kappa'$.

Case E-IFFALSE:

$(\text{cond false } e_2 e_3 \mid \nu \mid t) \rightarrow (e_3 \mid \nu \mid t + 1)$

We consider the typing derivation of $e \mid \nu$ in which we have $\Gamma \vdash e_3 : \tau \ \$ \ c_3, \kappa_3$ (from the premise T-COND) and $\nu : \Gamma \ \$ \ c'', \kappa''$ (from the premise T-COMPLETE). We combine them in a new application of T-COMPLETE to verify $\emptyset \vdash e_3 \mid \nu : \tau \ \$ \ c', \kappa'$.

Case E-IF:

$(\text{cond } e_1 e_2 e_3 \mid \nu \mid t) \rightarrow (\text{cond } e'_1 e_2 e_3 \mid \nu' \mid t + 1)$

By the premise of T-COND $e_1 : Bool$ and by L-COND $e_1 \mid \nu : Bool$. By applying the Inductive Hypothesis to a judgment with this as its last term, we obtain $e'_1 \mid \nu' : Bool$.

Looking at the typing derivation of $e'_1 \mid \nu' : Bool$:

$$g = \frac{\Gamma \vdash \nu(x_j) : \Gamma(x_j) \ \$ \ c_j, \kappa_j \text{ (for every } x_j \in \{x_1, \dots, x_n\} = \text{Domain}(\nu))}{\Gamma \vdash e'_1 : Bool \ \$ \ c_1, \kappa_1} \frac{h = (\emptyset \vdash \nu' : \Gamma \ \$ \ c'', \kappa'')}{\emptyset \vdash e'_1 \mid \nu' : Bool \ \$ \ c'_1, \kappa'_1} \text{(T-COMPLETE)} \text{(T-NU)}$$

Now we apply T-COND to g and the terms from the derivation of $\text{cond } e_1 e_2 e_3 \mid \nu$ (if needed, see α in L-COND for clarity):

$$i = \frac{\Gamma \vdash e'_1 : Bool \ \$ \ c_1, \kappa_1 \quad \Gamma \vdash e_2 : \tau \ \$ \ c_2, \kappa_2 \quad \Gamma \vdash e_3 : \tau \ \$ \ c_3, \kappa_3}{\Gamma \vdash \text{cond } e'_1 e_2 e_3 : \tau \ \$ \ c'_4, \kappa'_4} \text{(T-COND)}$$

Now we apply T-COMPLETE to h and i to verify: $\emptyset \vdash \text{cond } e'_1 e_2 e_3 \mid \nu' : \tau \ \$ \ c', \kappa'$

Case $e = x_i$ with $x_i : \tau_i$ and $x_i \in \text{Domain}(\nu)$ (as $e \mid \nu$ is closed):

By the definition of ν (and closure of the pair $e \mid \nu$) we rewrite $\vdash x_i \mid \nu : \tau \ \$ \ c', \kappa'$ as $\vdash x_i \mid \{x_i \mapsto \nu(x_i)\} \cup \nu_+ : \tau \ \$ \ c', \kappa'$. We now consider the possible evaluation cases individually:

Case E-VAR1 ($\nu(x_i) = v_i$):

$(x_i \mid \{x_i \mapsto v_i\} \cup \nu_+ \mid t) \rightarrow (v_i \mid \{x_i \mapsto v_i\} \cup \nu_+ \mid t + 1)$

We know from the derivation of $e \mid \nu$, the premise of T-NU gives $\Gamma \vdash \nu(x_i) : \tau_i \ \$ \ c_i, \kappa_i$. As $\nu(x_i) = v_i$ and $\Gamma(x_i) = \tau_i = \tau$ we can re-write this as $g = (\Gamma \vdash v_i : \tau \ \$ \ c_i, \kappa_i)$.

Again from the derivation we also have the premise of T-COMPLETE: $h = (\emptyset \vdash \nu : \Gamma \ \$ \ c'', \kappa'')$. Combining g and h under T-COMPLETE we can verify $\vdash v_i \mid \nu : \tau \ \$ \ c', \kappa' = \vdash e' \mid \nu : \tau \ \$ \ c', \kappa'$.

Case E-VAR2 ($\nu(x_i) = e_i$):

$(x_i \mid \{x_i \mapsto e_i\} \cup \nu_+ \mid t) \rightarrow (x_i \mid \{x_i \mapsto e'_i\} \cup \nu'_+ \mid t + 1)$

L-VAR gives us $(\emptyset \vdash \nu(x_i) \mid \nu : \tau \ \$ \ c', \kappa')$, which we rewrite by the definition of ν as $(\emptyset \vdash e_i \mid \nu : \tau \ \$ \ c', \kappa')$. Applying the Inductive Hypothesis we have that $(\emptyset \vdash e'_i \mid \nu' : \tau \ \$ \ c''', \kappa''')$. From the typing derivation of this expression we have a right side premise for T-COMPLETE: $h = (\emptyset \vdash \nu' : \Gamma \ \$ \ c'', \kappa'')$. Combining h with our prior left hand premise of T-COMPLETE, the T-VAR's $\Gamma \vdash x_i : \tau \ \$ \ 0, \emptyset$, we have:

$\vdash x_i \mid \nu' : \tau \ \$ \ c', \kappa' = \vdash e' \mid \nu' : \tau \ \$ \ c', \kappa'$.

Case $\mathbf{e} = \mathbf{let} \{x_1 = \mathbf{e}_1, \dots, x_n = \mathbf{e}_n\} \mathbf{in} \mathbf{e}_{n+1}$:

Case E-LETN:

$$(\mathbf{let} \{x_1 = \mathbf{e}_1, \dots, x_n = \mathbf{e}_n\} \mathbf{in} \mathbf{e}_{n+1} \mid \nu \mid t) \rightarrow (\mathbf{e}_{n+1} \mid \nu \cup \{x_1 \mapsto \mathbf{e}_1, \dots, x_n \mapsto \mathbf{e}_n\} \mid t + 1)$$

Thus our goal is to type: $\mathbf{e}_{n+1} \mid \nu \cup \{x_1 \mapsto \mathbf{e}_1, \dots, x_n \mapsto \mathbf{e}_n\}$

We first refer to the typing derivation of $\mathbf{e} \mid \nu$ (and define aliases to save space):

$$g = (\Gamma \vdash \mathbf{e}_1 : \tau_1 \ \$ \ c_1, \kappa_1 \ \dots \ \Gamma \vdash \mathbf{e}_n : \tau_n \ \$ \ c_n, \kappa_n)$$

$$h = (\Gamma, x_1 : \tau_1, \dots, x_n : \tau_n \vdash \mathbf{e}_{n+1} : \tau_{n+1} \ \$ \ c_{n+1}, \kappa_{n+1})$$

$$\mathcal{Y} = \frac{g \quad h}{\Gamma \vdash \mathbf{let} \{x_1 = \mathbf{e}_1, \dots, x_n = \mathbf{e}_n\} \mathbf{in} \mathbf{e}_{n+1} : \tau_{n+1} \ \$ \ c_{n+1} + \sum_{i=1}^n c_i, \kappa_{n+1} \cup \bigcup_{i=1}^n \kappa_i} \text{(T-LETN)}$$

$$\frac{\mathcal{Y} \quad \frac{i = (\Gamma \vdash \nu(x_j) : \Gamma(x_j) \ \$ \ c_j, \kappa_j \ \text{for every } x_j \in \{x_1, \dots, x_n\} = \mathit{Domain}(\nu))}{\emptyset \vdash \nu : \Gamma \ \$ \ c'', \kappa''} \text{(T-NU)}}{\emptyset \vdash \mathbf{e} \mid \nu : \tau \ \$ \ c_1, \kappa_1} \text{(T-COMPLETE)}$$

We call $\nu_1 = \{x_1 \mapsto \mathbf{e}_1, \dots, x_n \mapsto \mathbf{e}_n\}$ and call to g and h show that each $\nu_1(x_i) : \Gamma(x_i)$ for $x_i \in \mathit{Domain}(\nu_1)$. We see $\nu' = \nu \cup \nu_1$ and now we have what we need to apply T-COMPLETE and show our desired result.

$$\frac{h \quad \frac{\Gamma \vdash \nu'(x_j) : \Gamma(x_j) \ \$ \ c_j, \kappa_j \ \text{for every } x_j \in \{x_1, \dots, x_n\} = \mathit{Domain}(\nu')}{\emptyset \vdash \nu' : \Gamma \ \$ \ c'', \kappa''} \text{(T-NU)}}{\emptyset \vdash \mathbf{e}_{n+1} \mid \nu' : \tau \ \$ \ c', \kappa'} \text{(T-COMPLETE)}$$

Case $\mathbf{e} = \mathbf{trigger} \mathbf{e}_1 \ \mathbf{e}_2$:

Case E-TRIGGER1:

$$(\mathbf{trigger} \ \mathbf{e}_1 \ \mathbf{e}_2 \mid \nu \mid t) \rightarrow (\mathbf{cond} \ \mathbf{e}_1 \ \mathbf{e}_2 \ (\mathbf{trigger} \ \mathbf{e}_1 \ \mathbf{e}_2) \mid \nu \mid t + 1)$$

From the typing derivation of $\mathbf{e} \mid \nu$ we have $\mathbf{e}_1 : \mathit{Bool}$, $\mathbf{e}_2 : \tau$, $\mathbf{e} : \tau$, and $\emptyset \vdash \nu : \Gamma \ \$ \ c'', \kappa''$. We combine the first three under T-COND and the result with the last under T-COMPLETE to obtain our desired result.

Case E-TRIGGERDEG:

$$(\mathbf{trigger} \ \mathbf{v}_1 \ \mathbf{e}_2 \mid \nu \mid t) \rightarrow (\mathbf{cond} \ \mathbf{v}_1 \ \mathbf{e}_2 \ (\mathbf{trigger} \ \mathbf{e}_1 \ \mathbf{e}_2) \mid \nu \mid t + 1)$$

Identical to the above.

Case $\mathbf{e} = \mathbf{trigger} \{x_1 = \mathbf{e}_1, \dots, x_n = \mathbf{e}_n\} \mathbf{e}_{n+1} \ \mathbf{e}_{n+2}$:

Case E-TRIGGERLET:

$$(\mathbf{trigger} \ \{x_1 = \mathbf{e}_1, \dots, x_n = \mathbf{e}_n\} \ \mathbf{e}_{n+1} \ \mathbf{e}_{n+2} \mid \nu \mid t) \rightarrow (\mathbf{let} \ \{x_1 = \mathbf{e}_1, \dots, x_n = \mathbf{e}_n\} \ \mathbf{in} \ (\mathbf{cond} \ \mathbf{e}_{n+1} \ \mathbf{e}_{n+2} \ \mathbf{e}) \mid \nu \mid t + 1)$$

From the typing derivation of $\mathbf{e} \mid \nu$ we have $\mathbf{e}_{n+1} : \mathit{Bool}$, $\mathbf{e}_{n+2} : \tau$, and $\mathbf{e} : \tau$. We combine these under T-COND, the result with the variable premises under T-LETN and finally this with $\emptyset \vdash \nu : \Gamma \ \$ \ c'', \kappa''$ under T-COMPLETE to obtain our desired result.

Case $e = op\ e_1$:

Case E-OPAPPLY:

$(op\ v_1 \mid \nu \mid t) \rightarrow (v_2 \mid \nu \mid t + 1)$ The function *Apply* by definition, returns a value (which is inherently typed). By the typing derivation of $op\ v_1 \mid \nu$ we find $\emptyset \vdash \nu : \Gamma \$ c'', \kappa''$ and can apply T-COMPLETE to obtain our desired result.

Case E-OP1:

$(op\ e_1 \mid \nu \mid t) \rightarrow (op\ e'_1 \mid \nu' \mid t + 1)$

By the premise of T-OPAPPLY $e_1 : \tau_1$ and by L-OPAPPLY $e_1 \mid \nu : \tau_1$. By the Inductive Hypothesis we have $e'_1 \mid \nu' : \tau_1$.

Looking at the typing derivation of $e'_1 \mid \nu' : \tau_1$:

$$\frac{g = (\Gamma \vdash e'_1 : \tau_1 \$ c_1, \kappa_1) \quad h = (\emptyset \vdash \nu' : \Gamma \$ c'', \kappa'')}{\emptyset \vdash e'_1 \mid \nu' : \tau_1 \$ c', \kappa'} \text{(T-COMPLETE)}$$

Now we apply T-OPAPPLY to g and the terms from the derivation of $op\ e_1 \mid \nu$ (if needed, see α in L-OPAPPLY for clarity):

$$\frac{\Gamma \vdash op : \tau_1 \rightarrow \tau_2 \$ c_1, \kappa_1 \quad \Gamma \vdash e'_1 : \tau_1 \$ c_2, \kappa_2}{i = (\Gamma \vdash op\ e'_1 : \tau_2 \$ c', \kappa')} \text{(T-OPAPPLY)}$$

Now we apply T-COMPLETE to h and i to verify: $\emptyset \vdash op\ e'_1 \mid \nu' : \tau \$ c''', \kappa'''$

Case $e = get\ e_1$:

Case E-GETAPPLY:

$(get\ v_1 \mid \nu \mid t) \rightarrow (v_2 \mid \nu \mid t + 1)$

The function *get* by definition, returns a value (which is inherently typed). By the typing derivation of $get\ v_1 \mid \nu$ we find $\emptyset \vdash \nu : \Gamma \$ c'', \kappa''$ and can apply T-COMPLETE to obtain our desired result.

Case E-GET1:

$(get\ e_1 \mid \nu \mid t) \rightarrow (get\ e'_1 \mid \nu' \mid t + 1)$

By the premise of T-SENSORREAD $e_1 : Sensor\ \tau_1$ and by L-GET $e_1 \mid \nu : Sensor\ \tau_1$.

By the Inductive Hypothesis we have $e'_1 \mid \nu' : Sensor\ \tau_1$.

Looking at the typing derivation of $e'_1 \mid \nu' : Sensor\ \tau_1$:

$$\frac{g = (\Gamma \vdash e'_1 : Sensor\ \tau_1 \$ c_1, \kappa_1) \quad h = (\emptyset \vdash \nu' : \Gamma \$ c'', \kappa'')}{\emptyset \vdash e'_1 \mid \nu' : Sensor\ \tau_1 \$ c', \kappa'} \text{(T-COMPLETE)}$$

Now we apply T-SENSORREAD to g and the terms from the derivation of $get\ e_1 \mid \nu$ (if needed, see α in L-GET for clarity):

$$\frac{\Gamma \vdash get : Sensor\ \tau_1 \rightarrow \tau_1 \$ c_1, \kappa_1 \quad \Gamma \vdash e'_1 : Sensor\ \tau_1 \$ c_2, \kappa_2}{i = (\Gamma \vdash get\ e'_1 : \tau_1 \$ c', \kappa')} \text{(T-SENSORREAD)}$$

Now we apply T-COMPLETE to h and i to verify: $\emptyset \vdash get\ e'_1 \mid \nu' : \tau_1 \$ c''', \kappa'''$

4 Applications of the formalism

In this section we present concrete instantiations of the formalism presented in the previous section, in order to illustrate its benefit.

4.1 Additional Syntax

In addition to defining specific primitive operators (Opcodes) we also define the notion of a pair to allow these operators to accept multiple inputs.

e ::=	<i>expressions</i>
{e,e}	<i>pair</i>
ft	<i>flowtypes</i>
v ::=	<i>values</i>
{v,v}	<i>pair</i>
op ::=	<i>opcodes</i>
fst	<i>first projection of a pair</i>
snd	<i>second projection of a pair</i>
sizedimage	<i>allocate an image sensor (supporting specific sizes)</i>
add	<i>add integers</i>
facet	<i>count faces in an image</i>
resample	<i>change the resolution of an image</i>
ft ::=	<i>flowtypes</i>
deadline n e	<i>define a timing constraint</i>
τ ::=	<i>types</i>
$\tau \times \tau$	<i>pair</i>

4.2 Typing and Evaluation Rules for Pairs

Evaluation (strict) and typing rules for pairs are handled in a standard way.

$$(T\text{-PAIRCONS}) \frac{\Gamma \vdash \mathbf{e}_1 : \tau_1 \ \$ \ c_1, \kappa_1 \quad \Gamma \vdash \mathbf{e}_2 : \tau_2 \ \$ \ c_2, \kappa_2}{\Gamma \vdash \{\mathbf{e}_1, \mathbf{e}_2\} : \tau_1 \times \tau_2 \ \$ \ 1 + c_1 + c_2, \kappa_1 \cup \kappa_2}$$

$$(T\text{-PAIRFIRST}) \frac{\Gamma \vdash \mathbf{fst} : \tau_1 \times \tau_2 \rightarrow \tau_1 \ \$ \ 0, \emptyset \quad \Gamma \vdash \mathbf{e} : \tau_1 \times \tau_2 \ \$ \ c, \kappa}{\Gamma \vdash \mathbf{fst} \ \mathbf{e} : \tau_1 \ \$ \ 1 + c, \kappa}$$

$$(T\text{-PAIRSECOND}) \frac{\Gamma \vdash \mathbf{snd} : \tau_1 \times \tau_2 \rightarrow \tau_2 \ \$ \ 0, \emptyset \quad \Gamma \vdash \mathbf{e} : \tau_1 \times \tau_2 \ \$ \ c, \kappa}{\Gamma \vdash \mathbf{snd} \ \mathbf{e} : \tau_2 \ \$ \ 1 + c, \kappa}$$

$$(E\text{-PAIRFST}) \ \mathbf{fst} \ \{v_1, v_2\} \mid \nu \mid t \rightarrow v_1 \mid \nu \mid t \quad (E\text{-PAIRSND}) \ \mathbf{snd} \ \{v_1, v_2\} \mid \nu \mid t \rightarrow v_2 \mid \nu \mid t$$

$$(E\text{-PAIR1}) \frac{\mathbf{e}_1 \mid \nu \mid t \rightarrow \mathbf{e}'_1 \mid \nu' \mid t + 1}{\{\mathbf{e}_1, \mathbf{e}_2\} \mid \nu \mid t \rightarrow \{\mathbf{e}'_1, \mathbf{e}_2\} \mid \nu' \mid t + 1} \quad (E\text{-PAIR2}) \frac{\mathbf{e}_2 \mid \nu \mid t \rightarrow \mathbf{e}'_2 \mid \nu' \mid t + 1}{\{v_1, \mathbf{e}_2\} \mid \nu \mid t \rightarrow \{v_1, \mathbf{e}'_2\} \mid \nu' \mid t + 1}$$

4.3 Additional Typing Rules for Image Manipulation

Operations that manipulate images may explicitly specify a range of valid input sizes (*i.e.*, size constraints) on their input to ensure correct processing. In the example below, the face count opcode requires that its input be in the size range of 320 to 1024 (using manageable image widths as a size rather than actual total numbers of pixels which fall in the millions of pixels). The definitions of these opcodes implicitly combine an aspect of T-WEAKEN, by using size ranges (rather than single values) in their constrained size variables. This alternate approach has a distinct advantage over using T-WEAKEN prior to T-OPAPPLY (as is done in the Conditional in Section 3.5.3); namely the size and cost bounds are more accurate if they are computed against the input's actual size bound instead of weakened size bounds. It should be obvious to the reader that the *latentcost* values for the functions in this section have also been contrived in a manner to ease the presentation and discussion.

$$\begin{aligned}
 \text{(T-FACECT)} \quad & \frac{\text{type}(\text{facect}) = \text{Img}^{\{r_1, r_2\}} \rightarrow \text{Int}^{\{n_1, n_2\}} \quad \text{constr}(\text{facect}) = \kappa_1}{\text{facect} : \text{Img}^{\{r_1, r_2\}} \rightarrow \text{Int}^{\{n_1, n_2\}} \text{ \$ } 0, \kappa_1} \\
 & \kappa_1 = \{r_1 \geq 320, r_2 \leq 1024\} \\
 & \text{latentcost}(\text{facect}, \tau) = (\text{maxsize}(\tau)/2)
 \end{aligned}$$

A resampling operation is analogous to casting an image to be of a different size.⁵ It has no explicit values for the size variables of its input. Its only constraint is that the input be of some positive size (greater than zero).

$$\begin{aligned}
 \text{(T-RESAMPLE)} \quad & \frac{\text{type}(\text{resample}) = \text{Int}^{\{n_1, n_2\}} \times \text{Img}^{\{r_1, r_2\}} \rightarrow \text{Img}^{\{r_3, r_4\}} \quad \text{constr}(\text{resample}) = \kappa_1}{\text{resample} : \text{Int}^{\{n_1, n_2\}} \times \text{Img}^{\{r_1, r_2\}} \rightarrow \text{Img}^{\{r_3, r_4\}} \text{ \$ } 0, \kappa_1} \\
 & \kappa_1 = \{n_1 > 0, r_1 > 0, r_3 = r_1 * n_1, r_4 = r_2 * n_2\} \\
 & \text{latentcost}(\text{resample}, \tau) = (\text{maxsize}(\tau)/8)
 \end{aligned}$$

Image sensing hardware (*e.g.*, a web camera) has the capability to capture images in a range of possible resolutions. While the sized annotation makes sizing explicit in the type system, our programming language lacks explicit type annotation within its syntax. Thus, we formally define a new primitive to allocate a new image sensor (*ala* image) that accepts a resolution range explicitly, as an argument.

$$\begin{aligned}
 \text{(T-SIZEDIMAGE)} \quad & \frac{\text{type}(\text{sizedimage}) = \tau \rightarrow \text{Sensor } \text{Img}^{\{r_1, r_2\}} \quad \text{constr}(\text{sizedimage}) = \kappa_1}{\text{sizedimage} : \tau \rightarrow \text{Sensor } \text{Img}^{\{r_1, r_2\}} \text{ \$ } 0, \kappa_1} \\
 & \tau = \text{Int}^{\{n_1, n_2\}} \times \text{Int}^{\{n_3, n_4\}} \\
 & \kappa_1 = \{n_1 > 0, n_3 > 0, r_1 \geq \text{min}(n_1, n_3), r_2 \leq \text{max}(n_2, n_4)\} \\
 & \text{latentcost}(\text{sizedimage}, \tau) = (\text{maxsize}(\tau)/8)
 \end{aligned}$$

⁵While a resampling operation does increase the number of pixels in an image, it does not improve image *quality*. We will return to this issue in Section 5.

4.4 Flowtypes

Finally we introduce a new function whose sole purpose is to inject run-time constraints (a Flowtype in our nomenclature). This function annotates that its argument has an explicit deadline within the type system. For example, the example with function `period()` from the Introduction would be implemented as syntactic sugar using `deadline()`.

$$(T\text{-DEADLINE}) \frac{\Gamma \vdash n : \text{Int}^{\{n,n\}} \$ 0, \kappa_1 \quad \Gamma \vdash e_2 : \tau \$ c_2, \kappa_2}{\Gamma \vdash \text{deadline } n \ e_2 : \tau \$ c_2, \kappa_1 \cup \kappa_2 \cup \kappa_3}$$

where $\kappa_3 = \{c_2 \leq n\}$

$$(E\text{-DEADLINE}) \ (\text{deadline } n \ e_2 \mid \nu \mid t) \rightarrow (e_2 \mid \nu \mid t)$$

4.5 In Practice

In this section we have expanded our core language as much as possible so as to reflect our real operating and tasking environment. We will now construct several examples on which we apply our type system.

Inferring Optimal Image Resolution

In practice, a user may not provide a specific resolution (size) bound for images that are used as part of a larger computation in which the image is not part of the desired output. For example, a user who wishes to determine whether or not a light is on in a particular office is interested in a boolean result, not the intermediate image used to generate this result. As image sensors are able to capture images in a range of possible resolutions, our type system can use its size constraint system to suggest an optimal resolution, a range of feasible resolutions, or indicate that there is no feasible solution for the program as specified.

Example 1:

Face counting within an image from an image sensor with no explicit size bounds: `facect(get(image))`

$$\mathcal{D}_0 = \frac{\text{type}(\text{get}) = \text{Sensor } \tau \rightarrow \tau \quad \text{image} : \text{Sensor } \text{Img}^{\{r_1, r_2\}} \$ 0, \kappa_1}{\text{get}(\text{image}) : \text{Img}^{\{r_1, r_2\}} \$ c_1, \kappa_1} \begin{array}{l} \text{(T-IMAGESENSOR)} \\ \text{(T-SENSORREAD)} \end{array}$$

$$\mathcal{D}_1 = \frac{\text{type}(\text{facect}) = \text{Img}^{\{r_1, r_2\}} \rightarrow \text{Int}^{\{n_1, n_2\}} \quad \text{constr}(\text{facect}) = \kappa_2}{\text{facect} : \text{Img}^{\{r_1, r_2\}} \rightarrow \text{Int}^{\{n_1, n_2\}} \$ 0, \kappa_2} \text{(T-FACECT)}$$

$$\frac{\text{facect}(\text{get}(\text{image})) : \text{Int}^{\{n_1, n_2\}} \$ c_2, \kappa_1 \cup \kappa_2}{\mathcal{D}_1} \text{(T-OPAPPLY)}$$

$$\begin{aligned}
c_1 &= 1 + \mathit{latentcost}(\mathit{get}, r_2) = 1 + (r_2/8) \\
c_2 &= 1 + c_1 + \mathit{latentcost}(\mathit{facect}, r_2) = 1 + c_1 + r_2/2 \\
\kappa_1 &= \{r_1 \geq \mathbf{r}_{min}, r_2 \leq \mathbf{r}_{max}\} \\
\kappa_2 &= \{r_1 \geq 320, r_2 \leq 1024\}
\end{aligned}$$

Solving the constraints for r_1 and r_2 (we minimize r_1 and maximize r_2 subject to the constraints given above, we determine the simple constraints that the resolution of the image to manipulate, and thus the sensor itself, fall in the range $[320, 1024]$. Thus the SSD may allocate any available sensor that can produce images within this range of resolutions. A camera that can capture images at resolutions ranging from 1024 to 4096 is valid for this service fragment (provided it samples at 1024), as is a camera that can capture images at the range from 320 to 512.

Bear in mind that the minimum resolution in a range of resolutions is not always the most desirable value; while the minimum will consume the least computational resources, it may do so at the expense of computation confidence (*e.g.*, it may not be possible to detect all the faces in an image if the resolution is too small). As a result we advocate the use of the maximum size, provided there are resources available to accommodate the additional processing overhead. This processing overhead is easily measured by the cost function (c_2 , in the above).

Example 2:

Face counting within an image from an image sensor with no explicit size bounds *with an explicit deadline*: `deadline 322 facect (get(image))`

$$\frac{\frac{}{322 : \mathit{Int} \{n, n\} \$ 0, \kappa_0} \text{(T-INT)}}{\mathit{deadline} \ 322 \ \mathit{facect}(\mathit{get}(\mathit{image})) : \mathit{Int} \{n_1, n_2\} \$ c_2, \kappa_0 \cup \kappa_1 \cup \kappa_2 \cup \kappa_3} \mathcal{D}_1 \text{(T-DEADLINE)}$$

$$\begin{aligned}
c_1 &= 1 + \mathit{latentcost}(\mathit{get}, r_2) = 1 + (r_2/8) \\
c_2 &= 1 + c_1 + \mathit{latentcost}(\mathit{facect}, r_2) = 1 + c_1 + r_2/2 \\
\kappa_1 &= \{r_1 \geq \mathbf{r}_{min}, r_2 \leq \mathbf{r}_{max}\} \\
\kappa_2 &= \{r_1 \geq 320, r_2 \leq 1024\} \\
\kappa_0 &= \{n = 322\} \\
\kappa_3 &= \{c_2 \leq 322\}
\end{aligned}$$

In this example the valid range for capture is no longer $[320, 1024]$, when we solve for r_1 and r_2 we are limited to the range $[320, 512]$ as larger values of r_2 would exceed the constraint imposed by κ_3 .

It is worth noting that, despite slight differences in the typing derivation, this example has a size bound result that is identical to an example with explicit (user-specified) size bounds that are the same as those imposed by the `facect` opcode (`deadline 322 facect(get(sizedimage({320, 1024})))`).

Example 3:

Face counting within an image from an image sensor with no explicit size bounds *as the response in a trigger expression*: `trigger e0 facect(get(image))`

We assume the presence of a Boolean typed expression e_0 with cost c_0 and constraint set κ_0 .

$$\frac{e_0 : Bool \ \$ \ c_0, \kappa_0 \quad \mathcal{D}_1}{\text{trigger } e_0 \ \text{facect}(\text{get}(\text{image})) : Int^{\{n_1, n_2\}} \ \$ \ c_3, \kappa_0 \cup \kappa_1 \cup \kappa_2} \text{(T-TRIG)}$$

$$c_1 = 1 + \text{latentcost}(\text{get}, r_2) = 1 + (r_2/8)$$

$$c_2 = 1 + c_1 + \text{latentcost}(\text{facect}, r_2) = 1 + c_1 + r_2/2$$

$$\kappa_1 = \{r_1 \geq r_{min}, r_2 \leq r_{max}\}$$

$$\kappa_2 = \{r_1 \geq 320, r_2 \leq 1024\}$$

$$c_3 = (\mathcal{T} * c_0) + c_2$$

If expression e_0 starts running at time t_{start} , transitions to `true` at time t_{true} , and $\text{delay}(e_0)$ is the time to evaluate e_0 then we can write: $\mathcal{T} = \lceil \frac{t_{true} - t_{start}}{\text{delay}(e_0)} \rceil$. Recognizing that c_0 is a delay (bound) for e_0 , we can express this as $\mathcal{T} = \lceil \frac{t_{true} - c_0}{c_0} \rceil$. This example underscores the correctly bounding (or estimating) the “weight” of \mathcal{T} is dependent on estimating the factors required to cause e_0 to transition to true, which may include information about the sensing environment \mathcal{E} and costs determined elsewhere in the derivation (*i.e.*, to supply t_{start}).

Example 4:

Face counting within a *resampled* image, originally captured from an image sensor with no explicit size bounds: `facect (resample {4, get(image) })`

$$\mathcal{D}_2 = \frac{4 : Int^{\{n, n\}} \ \$0, \{n = 4\} \quad \mathcal{D}_0}{\{4, \text{get}(\text{image})\} : Int^{\{n, n\}} \times \text{Img}^{\{r_1, r_2\}} \ \$ \ c_1, \kappa_1 \cup \{n = 4\}} \text{(T-PAIRCONS)}$$

$$\mathcal{D}_3 = \frac{\text{type}(\text{resample}) = Int^{\{n, n\}} \times \text{Img}^{\{r_1, r_2\}} \rightarrow \text{Img}^{\{r_3, r_4\}} \quad \text{constr}(\text{resample}) = \kappa_2}{\text{resample} : Int^{\{n, n\}} \times \text{Img}^{\{r_1, r_2\}} \rightarrow \text{Img}^{\{r_3, r_4\}} \ \$ \ 0, \kappa_2} \text{(T-RESAMPLE)}$$

$$\mathcal{D}_4 = \frac{\mathcal{D}_2 \quad \mathcal{D}_3}{\text{resample} \{4, \text{get}(\text{image})\} : \text{Img}^{\{r_3, r_4\}} \ \$ \ c_2, \kappa_1 \cup \kappa_2 \cup \{n = 4\}} \text{(T-OPAPPLY)}$$

$$\frac{\text{type}(\text{facect}) = \text{Img}^{\{r_3, r_4\}} \rightarrow Int^{\{n_3, n_4\}} \quad \text{constr}(\text{facect}) = \kappa_1}{\text{facect} : \text{Img}^{\{r_3, r_4\}} \rightarrow Int^{\{n_3, n_4\}} \ \$ \ 0, \kappa_3} \text{(T-FACECT)}$$

$$\frac{\text{facect} : \text{Img}^{\{r_3, r_4\}} \rightarrow Int^{\{n_3, n_4\}} \ \$ \ 0, \kappa_3 \quad \mathcal{D}_4}{\text{facect} (\text{resample} \{4, \text{get}(\text{image})\}) : Int^{\{n_3, n_4\}} \ \$ \ c_3, \kappa_4} \text{(T-OPAPPLY)}$$

$$\begin{aligned}
c_1 &= 1 + \mathit{latentcost}(\mathit{get}, r_2) = 1 + (r_2/8) \\
c_2 &= \mathit{latentcost}(\mathit{resample}) + c_1 \\
c_3 &= \mathit{latentcost}(\mathit{facect}) + c_2 \\
\kappa_1 &= \{r_1 \geq \mathbf{r}_{min}, r_2 \leq \mathbf{r}_{max}\} \\
\kappa_2 &= \{n > 0, r_1 > 0, r_3 = r_1 * 4, r_4 = r_2 * 4\} \\
\kappa_3 &= \{r_3 \geq 320, r_4 \leq 1024\} \\
\kappa_4 &= \{n = 4\} \cup \kappa_1 \cup \kappa_2 \cup \kappa_3
\end{aligned}$$

We obtain that $r_1 * 4 \geq 320 \Rightarrow r_1 \geq 80$ and $r_2 * 4 \leq 1024 \Rightarrow r_2 \leq 256$. Minimizing for r_1 and maximizing for r_2 gives the resolution range $[80, 256]$ required of the sensor to be allocated for this fragment.

Worst-Case Computational Cost and Expression Size Bounds

The worst case computational cost is given as a result of our type system as the first argument after the $\$$ in a typing judgment. This is a scalar value and could be adjusted/calibrated as a function to estimate actual execution times on various physical resources. The worst-case computational bounds provided by our system provide one static-time validation mechanism to determine if explicit deadlines (or other run-time constraints given other Flowtypes) can not be met as specified. As presented, the `deadline` opcode relies on the unadjusted (discretized) computation bound.

In addition to static verification, the Service Dispatcher component of `SNBENCH` often must partition a task across physical resources if there is insufficient computing resources available on a single node to accommodate the entire task. The worst case computational cost, as presented, provides us with an initial metric to guide the allocation of physical computing nodes and sensory resources for a given task (and its constituent subtasks).

Similarly, for any given task or sub-task we can look at the upper bound size annotation on its type information to determine the potential network overhead associated with partitioning the larger task at that point.

4.6 Implementation Details

The type system described in this document has been implemented as part of our `SNBENCH` project. The implementation has been done in Java, making use of the open-source `JavaCC` project [Sri]. In its current incarnation it is presented as part of the `SNBENCH` development tool chain; the type checker is automatically invoked when compiling our high-level language (`SNAFU`) programming language to `STEP`. As the implementation is entirely modular and checks `STEP` code rather than `SNAFU` code, nothing stops us from using the type checker *apart from* the `SNAFU` compiler, however (1) `SNAFU` is eminently more readable than `STEP` and (2) the implementation includes several hooks to track line numbers within `SNAFU` code as an error tracking convenience to its users. Put differently, while the type system can exist without `SNAFU`, we have gone to some lengths to ensure `SNAFU` users can enjoy the benefit of this type checker as well.

Our implementation of the type checking engine includes a broader range of STEP constructs than the Core STEP that is presented here, however there are some limitations to this support; for constructs for which we lack the complete formalism use heuristics to bound cost and default back to basic, unsized type checking rather than annotated type checking for opcodes which lack size constraint annotation.

To solve the constraint sets that are built as a result of the sized type checking, we invoke the GNU Linear Programming Toolkit (GLPK) [And], which can be used to solve a system of constraints for linear programming, and mixed integer programming. The decision to use GLPK is based on project maturity, community support and API availability. At present we emit our sizing constraints in the GNU MathProg language and invoke the GLPK via external script, though nothing (other than time) precludes finer integration via the GLPK Java Interface [Bjo].

Figures 2 and 3 are screen shots taken from the the SNAFU development environment that feature results generated from the implementation of the type checker.

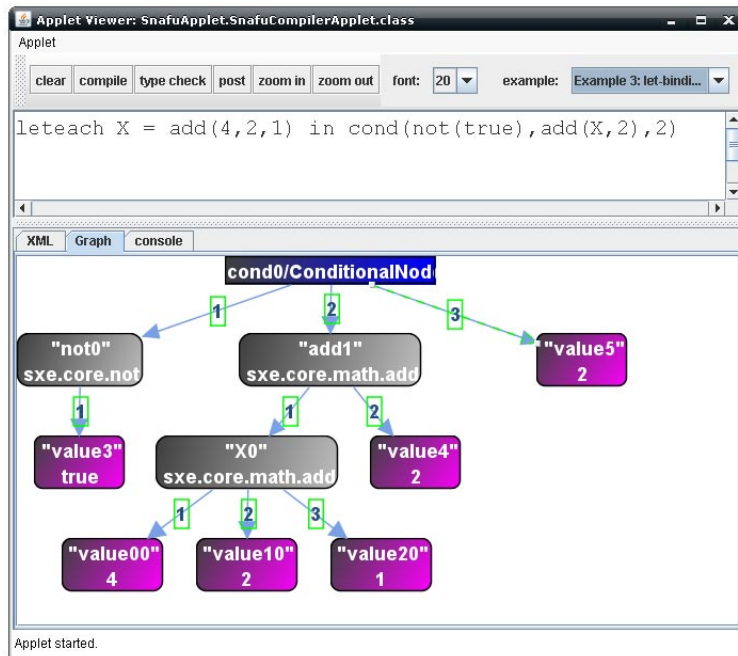


Figure 2: A screen shot from the SNAFU compiler showing the successful type checking of a STEP program. By default, the feedback is given graphically to the user.

5 Future Work

5.1 Additional Type Annotations

In this paper we have presented upper and lower size bound data type annotation, yet other useful annotations exist and can be easily integrated into this type system by extending the existing annotation pair to a tuple or larger ordered set and defining the desired subtyping relation.

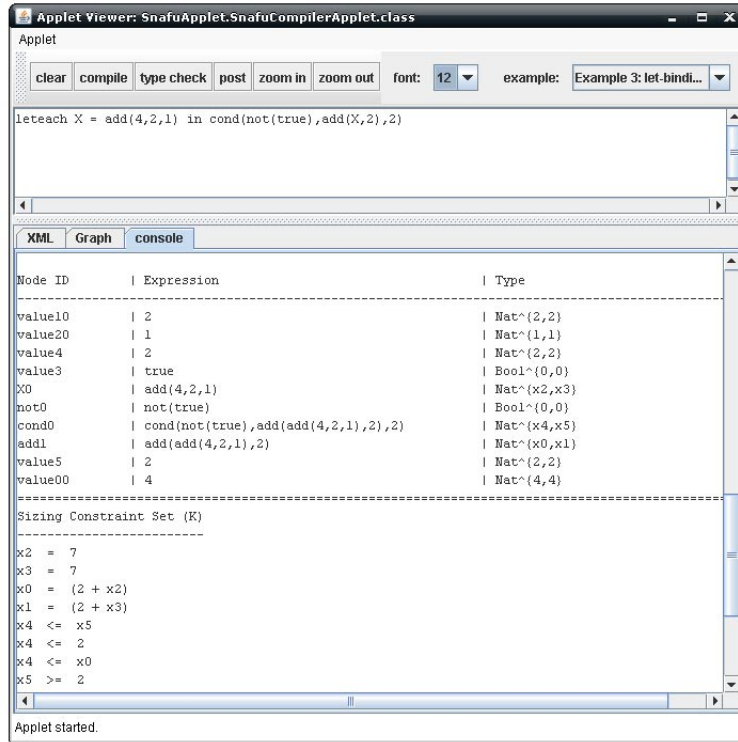


Figure 3: A screen shot from the SNAFU compiler showing size annotation constraints that result from type checking a STEP program.

One such example is the notion of image quality, which is different from data size. Data *size* is used bound the operational requirements of a function (including those that manipulate images), whereas image *quality* speaks to the valid data in the image. As far as operational/functional correctness is concerned a resized image is operationally valid, however with respect to the desired programmatic output, a smaller image that has been resized to a larger resolution is not truly interchangeable with an image captured at a larger resolution. When an image is resized (*e.g.*, via scaling or resampling, say) the size (resolution) of the image changes no longer reflects the number of data points it contained originally (we are calling this “quality”).

We could easily support a notion of an image’s quality within our type annotations, by adding a dimension for the “lowest” value that we have ever seen for an image’s lower size bound.⁶ This value could distinguish between a true high resolution image and data that has been up-cast or coerced to satisfy a function’s size constraints. Quality also has a well established meaning with respect to numerical data as well, and might be defined to reflect potential for rounding errors, data accuracy, *etc.* Certainly other image and video related aspects could be tracked as well, including color-depth for images, frame rates for video, and so on.

⁶Recall the definition of T-RESAMPLE increases both the upper and the lower size bound.

5.2 Applications to Image Pyramids

Image Pyramids [Kro91],[AH91] are a well established technique in the field of image processing. The technique involves maintaining multiple copies of the same image at different resolutions (the a hierarchy of resolutions form a logical pyramid) such that the appropriate resolution can be selected for manipulation depending on the needs of the manipulation function. The applications of this formalism to this community is potentially two-fold. (1) We can extend the type system to include an image pyramid as a first class type and extend the annotations to support the list of resolutions available in the pyramid. (2) Static analysis of the image processing flow can tell us precisely which resolutions need to be kept in the pyramid and which can be removed. In the latter case, the potential benefit of SNBENCH and sized typing is quite significant as it might be possible to remove the pyramid entirely. For distributed computations, we can split the pyramid into individual image instances based on the analysis of size/use and ensure that we are passing as few copies of the image (inside the pyramid) and as few copies of the pyramid itself, as possible.

5.3 Cost and Size Signatures For New Opcodes

The operational correctness of the type system is contingent on the presence of accurate size, cost and constraint data for Opcodes (primitive operators). SNBENCH provides a facility by which new Opcodes may be added to a service library quickly and easily, through the implementation of a simple Java interface. At present the type system maintains an embedded definition for the latent costs and size constraints of the current Opcode library, however for the sustainability of the type system, it is essential that these definitions are provided by the Opcode authors and automatically extracted directly from the Opcode definitions. Beside changing the Opcode implementation interface, the type system must also change to import the rules from the Opcode library directly. There is also a concern that Opcode authors might specify very weak size constraints and very high costs (possibly lacking the knowledge to do so correctly) and that the end result is a type system that is only as strong as its weakest link. Any future work that would automatically extract this information from an Opcode implementation would be a fantastic solution to this problem (and others).

5.4 Different Models of Cost

At present our type system provides a single model of computational cost, a worst-case upper bound. As there is the possibility for multiple size annotations (dimensions) of types, there is an opportunity to provide multiple cost metrics based on these other dimensions. Defining a so-called “minimum” computational cost (or optimal case) would be trivial, and other cost models including non-computation costs models (*e.g.*, defining some financial cost, or total memory utilization cost) would be possible, if not trivial. The approach that seems the most straight forward to enabling multiple cost models would be to provide functional costs, such that all costs computed on the right hand side of the \$ would be user/programmer customizable and configurable as a function of the data available in the typing rules.

5.5 Expanding to “Complete” STEP

At present we support only a subset of the STEP programming language. The remaining STEP constructs (*e.g.*, streaming/persistent triggers and their associated read semantics) are difficult to represent as a result of their asynchronous execution and complex internal state. Despite these challenges, we have worked to establish a typing formalism for these constructs however they are presently incomplete. Despite this, the present limited form is useful for non-persistent, “straight-shot” programs (which actually capture a significant amount of image processing tasks). For persistent/streaming trigger constructs we apply a simple heuristic: we solve the computable portion and multiply by estimates to derive the approximate costs for services that include non-Core portions of STEP. Regardless of this makeshift solution, it would be advantageous (and naturally appealing) to have a type system that includes the complete STEP language.

6 Conclusions

In this paper, we have presented our formal type system for multi-dimensional sized types. Unlike other sized type systems our work tracks both an upper and lower size bound for data, defines a logical subtype relation for images capable of bounding computation, maintaining functional correctness, and deduce feasible data sizes from implicit and explicit constraints within program fragments. We presented this system and have provided examples that illustrate the use of the type system. We are confident in the many potential uses for this formalism to the Image processing and Sense and Respond communities.

References

- [AH91] F. Ackermann and M. Hahn. Image pyramids for digital photogrammetry. *Digital Photogrammetric Systems*, pages 43–59, 1991.
- [And] Andrew Makhorin, Department for Applied Informatics, Moscow Aviation Institute, Moscow, Russia. GLPK (GNU Linear Programming Kit). <http://www.gnu.org/software/glpk/>.
- [BBKO05] Azer Bestavros, Adam Bradley, Assaf Kfoury, and Michael Ocean. SNBENCH: A Development and Run-Time Platform for Rapid Deployment of Sensor Network Applications. In *IEEE International Workshop on Broadband Advanced Sensor Networks (Basenets)*, October, 2005.
- [Bjo] Bjoern Frank. GLPK 4.8 Java Interface. <http://bjoern.dapnet.de/glpk/index.htm>.
- [Bos] Boston University, Department of Computer Science. Sensorium Research Homepage. <http://www.cs.bu.edu/groups/sensorium/>.
- [HFH06] Kevin Hammond, Christian Ferdinand, and Reinhold Heckmann. Towards formally verifiable resource bounds for real-time embedded systems. *SIGBED Rev.*, 3(4):27–36, 2006.
- [HPS96] John Hughes, Lars Pareto, and Amr Sabry. Proving the correctness of reactive systems using sized types. In *Symposium on Principles of Programming Languages*, pages 410–423, 1996.

- [Kro91] W. Kropatsch. Image pyramids and curves. an overview. Technical report, Department for Pattern Recognition and Image Processing of the Institute of Automation, University of Technology, Vienna, Austria, 1991.
- [LH96] Hans-Wolfgang Loidl and Kevin Hammond. A sized time system for a parallel functional language. In *Proceedings of the Glasgow Workshop on Functional Programming*, Ullapool, Scotland, July 1996.
- [OBK06] Michael J. Ocean, Azer Bestavros, and Assaf J. Kfoury. SNBENCH: Programming and Virtualization Framework for Distributed Multitasking Sensor Networks. In *VEE '06: Proceedings of the 2nd international conference on Virtual execution environments*, pages 89–99, New York, NY, USA, 2006. ACM Press.
- [RG94] Brian Reistad and David K. Gifford. Static dependent costs for estimating execution time. In *LISP and Functional Programming*, pages 65–78, 1994.
- [Sri] Sriram Sankar, Metameta. Java Compiler Compiler [tm] (JavaCC [tm]) - The Java Parser Generator. <https://javacc.dev.java.net/>.