

Procedural Graph-store Processing

Eric Griffis

University of California, Los Angeles
egriffis@cs.ucla.edu

Abstract

Graph-store Processing (GSP) is a novel programming paradigm for managing, sharing, and computing on arbitrarily structured and possibly distributed collections of digital information. As an informal introduction to basic GSP concepts, I present *pGrasp*, an experimental language for specifying computations as procedures within the context of a *graph-store*—a directed property graph over variable contexts. In this report, I outline the need for a paradigm like GSP. I introduce basic GSP concepts and relate them to *pGrasp* constructs. I produce a series of *pGrasp* programs which culminates in a simple distributed task queue with interesting properties. I provide a formal specification for *pGrasp*, along with an overview of a prototype implementation. Finally, I discuss in detail the relationship between *pGrasp* and GSP, and describe the next steps toward a practical GSP platform.

1. Introduction

We have, generally speaking, many tools at our disposal for solving a given problem in software. Despite the fact that all Turing-complete programming languages are equally expressive, each language tends to excel in some aspects at the expense others. Datalog, a logic programming language that models computations as queries over sets of related terms, solves problems in deductive reasoning elegantly but has some trouble with evolving state. C, an imperative language, handles mutable state well but can be difficult to reason about. Indeed, a language perfect for every situation would obviate the need for all others.

Thus, a useful programming language inevitably fills a niche—a setting in which particular kinds of problems can be solved *naturally*, with minimal complication. This reasoning applies to programming paradigms as well. Object-oriented programming (OOP) decomposes large problems into independent components to maximize productivity. Symbolic programming treats code as data to maximize flexibility. Imperative programming operates close to the underlying hardware platform to maximize performance.

How about for Web information systems [6], or other forms of distributed information systems? What kinds of processes can we execute on loosely structure information repositories like the Web, and how should we structure such efforts? What are the common properties of distributed, ad-hoc information

systems? Can we exploit these commonalities to reduce complexity in distributed systems? To summarize these issues, I pose the following question.

How can we most naturally manage, share, and compute on loosely structured digital information?

With respect to existing tools and methods, I see no clear answer to this question. For the Web, we might argue that the current trajectory of Web technologies is acceptable and will continue to improve with time. The evidence is, however, in direct opposition to this argument. The current state of Web software development is quite complex. Indeed, we must maintain an entire software “stack” of distinct products in order to enact a presence on the modern Web, and plans for the Semantic Web promise to inject an entirely new sub-stack into the existing model. Consider the volume of distinct configuration, description, and general programming languages involved in the following scenario.

To share information on the modern Web, we require at minimum a Web server and some HTML pages. To adjust elements of presentation, we use Cascading Stylesheets (CSS). We also use JavaScript for general computations like looping, interactive event handling, and back-end communications via AJAX libraries that transmit messages as JSON or XML documents. If we intend to offer proprietary or complex services, we integrate a server-side application development framework like PHP into our Web server. To manage structured data, we incorporate an SQL database or some form of NoSQL alternative. If we intend our service to become large, we might implement our applications on the SOAP messaging framework. If growth occurs unexpectedly, we might choose CORBA, an IDL-based middleware, instead.

No single technology in this stack is designed to simultaneously manage, share, and compute on a structured information repository. Moreover, each component presents a distinct perspective on the nature of the system as a whole. Disparate languages and priorities lead to duplicated effort and complicated work-arounds. Furthermore, the Web is but one form of distributed information system. We also have, for example, e-mail, video chat, and instant messaging systems, not to mention the plethora of emerging systems unique to mobile platforms. Each form poses different goals and challenges:

push versus pull, stream versus block, reliability versus performance, and so on.

Despite the profusion of concepts and languages throughout the stack, we can observe several common elements. Virtually all back-end processes communicate via tree structures, ranging in complexity from flat key-value headers to complete XML or JSON documents. Each process in the system encodes, decodes, consumes, transforms, and produces structured messages to affect its contribution to the overall behavior of the system. These messages represent structures local to each process, requests to act on these local structures, or responses to such requests.

There are clear opportunities to simplify the stack. Efficient message passing and processing are crucial elements of any distributed software system. Every layer of the stack stands to benefit from a natural setting in which to inspect and operate on messages directly, without spurious translations between representations. Furthermore, each implementation may depend on messaging internally, in the form of OOP method calls. To address this clear opportunity, I propose the graph-store.

1.1 Graph-store Processing

A graph-store is a graph-based data structure, designed specifically to manage, share, and compute on arbitrarily structured collections of information. The nodes of a graph-store are similar to C structs or OOP objects. A graph-store node differs from a struct in that it can be used directly as a context in which to evaluate a program. It differs from an object in that we can choose dynamically in which node to evaluate each sub-program. In other words, a graph-store node is like an object with methods that affect the internal state of other objects.

This arrangement flagrantly violates the principle of encapsulation. As a principled approach to software construction under such conditions, I propose the Graph-store Processing (GSP) paradigm. GSP appeals to graph rewriting [10, 11] as a basis for structuring and analyzing large programs with formal methods similar to functional programming, but in total absence of closed-world, top-down notions of modularity such as objects or functions—static barriers through which information passes according to code producers. Instead, I adopt an open-world, bottom-up approach, leveraging concepts such as functional pattern matching [7] and multi-staged programming (MSP) [13]—tools for code consumers to inspect and affect sub-program behavior dynamically. GSP stands apart from previous work with graph rewriting [1, 5] in its use of staging as the primary means of abstraction and its emphasis on simple sharing constructs.

GSP strikes a balance between the symbolic and functional paradigms. Specifically, GSP comprises a novel combination of staged evaluation and syntactic pattern matching. The result is a complete inversion of the concept of modularity. In

a modular system, we attach names to *code*—each component is a black box that distills arguments into return values. In a GSP system, we attach names to *data*—each component is an arbitrarily complex message, in the object-oriented sense [4], that describes a computation to be carried out through a series of rewrites.

In an imperative language like Java, the methods of an object describe *how* particular tasks are carried out as sequences of instructions that affect its internal state. Object boundaries insulate internal states so methods can not rely on transient external details or lead otherwise to inconsistencies. A GSP language specifies *what* a solution to a particular problem is in a declarative style. A GSP node consumes a sub-program, alters or acts on the parts that it understands, and passes the modified sub-program on as its result, possibly for further modification by other nodes—the sub-program *is* the argument, internal state, and return value.

This is not to say that modularity is impossible in a GSP system, just that modularity is extremely difficult to simulate within an individual graph-store and merely undesirable in general. In a distributed setting where local processes operate within distinct graph-stores, the graph-stores are de-facto modules, if only for the underlying physical boundaries. GSP defeats this limitation elegantly by prescribing a primitive node transmission mechanism.

My decision to abandon modularity might seem absurd at first, but the rewrite model of computation compensates amply by appealing to our extra-computational intuition; it is, for example, naturally concurrent. According to José Meseguer [8],

The idea of concurrent rewriting is very simple. It is the idea of *equational simplification* that we are all familiar with from our secondary school days, *plus* the obvious remark that we can do many of those simplifications independently, i.e., in *parallel*.

This excerpt highlights the peculiar nature of rewriting, about which reasoning leads typically from a simple premise directly to a desirable property. This quality allows us to express and compose a broad range of behaviors naturally in a bottom-up fashion, as the following comparison illustrates.

To implement matrices in a modular language without native support, we write a function library that constructs and operates on native representations of matrix data. Before we can use this library, we must understand its API. If the language contains meta-programming features, we can simplify the notation with additional work. Given two distinct matrix libraries, we either choose one statically or produce a compatibility layer that encapsulates our dynamic selection logic.

In a modular setting, we build programs by selecting implementations.

To implement matrices in a GSP language, we specify a set of rules that rewrite operations over matrices in an intuitive notation into, say, matrices of operations over scalars. We use the library by applying its rules repeatedly and recursively until no further rewrites are possible. Given an alternative implementation with identical notation, we choose one statically or dynamically and have no need for a compatibility layer. If the alternative follows a different notation, our compatibility layer contains only rules to map between notations.

In a graph-store setting, we build programs by selecting notations.

This report proceeds as follows. Section 2 describes the pGrasp programming language and its core constructs. Section 3 introduces pGrasp programming via three example programming sessions. From core constructs, I implement lists and queues. From these structures, I implement a networked queue management API for community resource pooling. Section 4 specifies the syntax and semantics of the pGrasp language formally, followed by details in section 5 of the prototype implementation that informed the formal design process. Section 6 briefly discusses discrepancies between pGrasp and GSP. Section 7 concludes with a summary of possible next steps for the GSP project.

2. The pGrasp Language

pGrasp is a procedural programming language designed to highlight some core principles of GSP, but the pGrasp language offers neither staged evaluation nor functional pattern matching. It is a proof of concept which favors familiar constructs over innovative design. I work around its deficiencies with convenience constructs and conventions borrowed from low-level languages. The language, nevertheless, is useful and admits compact solutions to real problems.

A pGrasp program constructs and manages a *graph-store*, a labeled, directed property graph that exhibits index-free adjacency of nodes, where each node is a distinct variable name space. pGrasp provides a variety of features, many of which are familiar. A pGrasp program is a semicolon-delimited sequence of statements. We can store and recall variables by name. We have looping and branching. We can capture and apply sub-programs as procedures, and interact with external resources via opaque descriptors. I use common data types like booleans, numbers, and strings in a conventional syntax. For slightly more detail, see the formal syntax specification in section 4.1.

Program evaluation progresses from left to right and top to bottom. I use parentheses (resp., braces) to denote explicit alternative groupings or orders of evaluation of expressions (resp., statements). Often, I group terms unnecessarily—especially for procedures, loops, and conditionals—to convey some intuition for constructs with similar behavior in other languages. Reserved words appear in a proportional sans-serif font, like this. Identifiers appear in a monospaced serif font, like here.


A variable, node reference, cursor (sec. 2.4), or value appears as the italicized letter x , n , c , or v , respectively, possibly with super- or sub-scripts.

2.1 Nodes and References


A graph-store instance contains a heap of nodes and a pair of special variables. A *heap* is an indexed set of nodes. A *node* is a dictionary, i.e., a set of mappings from variable names to values. To access a variable, we must obtain a reference to the node it inhabits. A *reference* is an opaque token that identifies a resource. We create a directed link between two nodes by setting a variable in the parent to a reference of the child. Occasionally, I conflate the notions of node, reference, and variable, assuming that the concept at hand is apparent.

A new pGrasp instance contains exactly one non-removable node called the *root*. Initially, both special variables reference the root. The first is immutable and therefore always references the root. The second references the *current node* and is always mutable. We can access the root and current node references with keywords `root` and `this`.

The following diagram shows the graph-store of a fresh pGrasp instance. In a graph-store diagram, each node appears as a box with rounded edges that contains its identifying reference at top, followed by any variables that do not contain a node reference. The special references n^r and n^c identify the root and current node, respectively. For each stored node reference, an edge labeled by the variable name points from parent to child.




The next diagram shows the heap of a fresh pGrasp instance. In a heap diagram, node references appear on the left. The variables within a node appear to the right of the corresponding reference as a chain of name-value pairs. An empty box represents a node with no variables.




We have several constructs available for creating new variables, but the most direct construct is the set statement. To set the root variable `toggle` to the boolean value `true`, write:

```
set root toggle true;
```


This statement grows the root node in the heap by one mapping.



Equivalently, the graph-store root gains a single non-reference mapping.




Graph-store diagrams are also useful for describing abstract processes as graph rewrite rules, as in the following illustration of the semantics for the set statement. For convenience, I assume that the left-hand side of a rule matches any node with at least the displayed elements, that the right-hand side of a rule replaces only the elements matched by the left-hand side, and that we apply rules repeatedly and recursively in any convenient order until we exhaust all possible matches.



We also have many ways to create nodes. The most basic, however, is the node statement. Because every non-root node must have a parent, all node construction statements require a parent node and a variable name in which to store a reference to the new node. Here is how we create a fresh node and store its reference in the root variable *n*.

```
node root n;
```


This statement adds a new node to the heap and grows the root node by one mapping.




Now that we have a non-root node in the heap, we can set the current node directly with the *in* statement.

```
in n;
```

This copies the reference stored in root variable *n* into the current node special variable.



Here are the rewrite rules for the node and in statements.



We have two constructs for removing data. To delete a single variable, use the *unset* statement. If the deleted variable contains the sole copy of a node reference, then the referenced node is removed from the heap. Removing a node deletes all of its variables implicitly, so a cascade effect is possible. We can also use the *reset* statement to delete all of the variables in a node at once, as in the following sequence which leverages the cascade effect to empty the graph-store completely.

```
in root; reset root;
```

2.2 Paths

We have two forms of variable resolution available, each corresponding to one of the special variables. Most simply, a lone variable name resolves to the value it maps to in the current node, or null if no such mapping exists. We can also access variables outside the current node via path expressions. A *path* is a period-delimited sequence of sub-expressions that each, except possibly the last, evaluate to a node reference.

pGrasp evaluates the first sub-expression of a path in the current node. Because pGrasp evaluates each remaining sub-expression in the node referenced by the previous, all but possibly the last must evaluate to a node reference. If a variable does not exist in the appropriate node, the result is the null value. In conditional test expressions, null and false are equivalent. To test for a null result directly, use an equality test expression; e.g., `x == null`. In its current form, pGrasp offers no mechanism to differentiate between a variable set to null and a variable that does not exist, hence setting a variable to null effectively unsets the variable.

Absolute paths begin with the root keyword.¹ A relative path is any path that is not an absolute path. The following statement uses an absolute path to set the current node to the node referenced by root variable *n*.

```
in root.n;
```

Because path sub-expressions are more-or-less arbitrary, we can embed branch selections directly into the path. The next example sets the current node to either `root.left.target` or `root.right.target`, depending on the value of variable `root.toggle`.

```
in root.(if toggle then left else right)
.target;
```

Observe that pGrasp does not resolve path sub-expression variables in the current node implicitly. If we want to access a current node variable from within a path sub-expression, we must specify so explicitly. For instance, the following statement conditions the path from the previous example on the current node variable `toggle`.

```
in root.(if this.toggle then left else right)
.target;
```

¹ To avoid clutter, I assume that the root variables are always in scope for the remainder of this document.


2.3 Procedures

Procedure semantics in pGrasp are extremely simple because the run-time environment does not maintain an implicit call stack. Instead, a called procedure stores its arguments in the current node and then replaces the calling statement with its body. By convention, the caller of a procedure must back up any important variables in the current node before issuing the call.

Procedure construction is simple. To create a procedure, use the `proc` statement. Take, for example, the following implementation of a list data structure.

```
proc root nil () {set this nil? true};
proc root cons (car,cdr) {set this pair? true};
```

These procedures are essentially data constructors—they embed a particular structure within the current node. The first procedure, `root.nil`, takes no arguments. When called, it sets the current node variable `nil?` to true. The second procedure, `root.cons`, takes two arguments: `car` and `cdr`. When called, it stores its arguments in the current node and then sets the current node variable `pair?` to true. The following rewrite rules illustrate these semantics.



For convenience, pGrasp provides four procedure calling constructs. The call statement is the fundamental procedure call construct; it simply installs the arguments into the current node and evaluates the body. The call statement is a useful tool for exploring alternatives to traditional stack-based procedure call semantics for, e.g., various forms of distributed task scheduling [2], parsing ambiguous grammars [12], or various other cases involving non-determinism that might be desirable in a distributed setting.

The build statement behaves like the call statement, but changes the current node variable for the duration of the call only. This is useful for directing simple constructor-like procedures, such as `nil` and `cons`, into an existing node. The following example constructs an empty list and stores a reference to it in the current node variable `L`.

```
node this L; build L nil ();
```

Because of the complexity involved in managing the call stack explicitly, the code and build statements are of limited utility. For higher-level programming tasks, the remaining two constructs provide the familiar procedure call semantics. The new statement behaves like the build statement, but first installs a fresh current node. The return statement behaves

similarly, but keeps only a single variable from the new node. In this next example, either statement is equivalent to the previous example.

```
new this L nil ();
return this L nil () this;
```

2.4 I/O

A graph-store instance contains, in addition to its heap and special variables, a communications dispatch object. A *dispatch object* is an indexed set of pairs of byte queues. A *cursor* is a reference to a byte queue pair. In order to connect to external resources like human input devices and network sockets, we submit a description of the desired resource to the dispatch object. If the resource is available, the dispatch object returns a cursor. When the external resource is another pGrasp instance, we can exchange sub-graph-stores with simple primitive constructs.

pGrasp mediates all external communications through its dispatch object. To connect to a resource, use the open statement. This construct takes a node describing the requested resource, along with a target node and variable name. The target variable receives a cursor representing a connection to the resource if successful, or null otherwise. For convenience, I assume the existence of two universal cursors, accessible by keywords `stdio` and `stderr`. No standard external resource request format exists yet, so I assume the following constructors produce valid requests for network sockets.

```
proc root tcp (host,port) {set this tcp? true};
proc root udp (host,port) {set this udp? true};
```

We can exchange bytes over an established connection with the send and rcv statements. To terminate an established connection, use the close statement. The following example creates a TCP connection to the UCLA public web server, sends a simple HTTP request, blocks until it receives the first TCP packet of the HTTP response, then terminates the connection.

```
new this addr tcp ("www.ucla.edu", 80);
open addr this ucla;
send ucla "GET / HTTP/1.0\r\n";
rcv ucla this response;
close ucla;
```

When connected to another pGrasp instance, we can exchange entire sub-graph-stores with the print and read statements. This simple mechanism extends the principles of GSP into the realm of distributed computing.

Suppose that pGrasp instance A can request a listening TCP socket on host/port pair `example.com/3636`, which will block until a connection is established. Suppose, further, that pGrasp instance B has a structured metadata collection rooted at `project.data`, which includes any procedures necessary for managing the collection. Then, the following example

transmits a complete copy of the collection and management code from A to B.

```

/* in instance A */
new this addr tcp ("example.com", 3636);
open addr this B;
print B project.data;
close B;

/* in instance B */
new this addr listen ("0.0.0.0", 3636);
open addr this A;
node this project;
read A project data;
close A;


```

These high-level constructs for structured data interchange reduce the barriers for entry to a broad range of distributed computing tasks. Virtually every communication layer of a modern web software stack fits naturally into this model. Furthermore, because procedures and evaluation contexts are just data, computations can be distributed easily across a network.

3. Example Sessions

3.1 Lists

Suppose we start a pGrasp session that contains only the aforementioned list constructors, as illustrated in the following heap diagram.




First, create a work space and set it as the current node.

```

node root work;
in work;

```

The heap now has the following structure.




Next, construct an empty list.

```

new this L nil ();

```

We now have a reference to an empty list in the current node variable L.



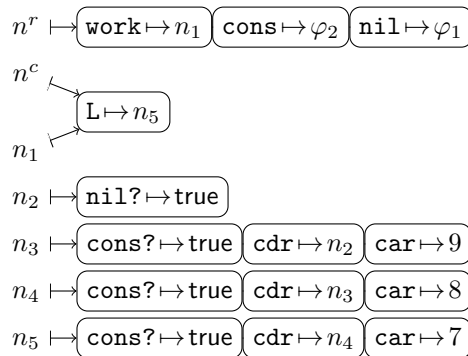
Extend the list by three numbers.

```

new this L cons (9,L);
new this L cons (8,L);
new this L cons (7,L);

```

Each new statement overwrites L with a link to a fresh cons cell. The list appears to grow because the cdr of each new link receives the old value of L.




We can calculate the length of L in an imperative style. First, initialize a counter and make a working copy of L.

```

set this N 0;
set this X L;

```

The current node now has the following structure.




Next, write the counting loop.

```


while (not X.nil?) {
  set this N (N + 1);
  set this X X.cdr }

```

Each iteration increments N by one and discards a single link. After the first iteration, X references the second element of L.




After the third iteration, X references the end of L and the loop terminates.



Delete X, as it is no longer useful.

```
unset this X;
```

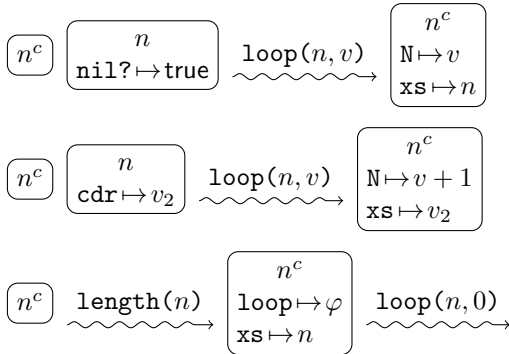
The current node appears as expected.



The following procedure calculates the length of a list, but in a functional style. In this case, the call statements implement tail recursion.

```
proc root length (xs) {
  proc this loop (xs,N) {
    when (not xs.nil?)
      call loop (xs.cdr,N+1);
  };
  call loop (xs,0) }
```


The rewrite rules for this procedure are not so obvious. Since the call statement essentially applies an arbitrary rewrite rule, I represent the loop initiation call as a dangling transition arrow. This notational kludge is an unfortunate consequence of the tension between modular abstraction (procedures) and staging (rewrite rules). In a truly non-modular setting, such artifacts do not arise.




Because pGrasp procedures do not return values in the traditional sense, each procedure must specify a convention for identifying its results. Though more sophisticated conventions are possible, I assume that the result is apparent by inspection. In this case, N contains the result. Since N contains the only interesting value, we use the return statement when calling this procedure.

```
return this M length (L) N;
```


The return statement sets the call up in a fresh working node, and the procedure body installs the inner loop immediately afterward. Then, the first call to loop sets up the counter N.



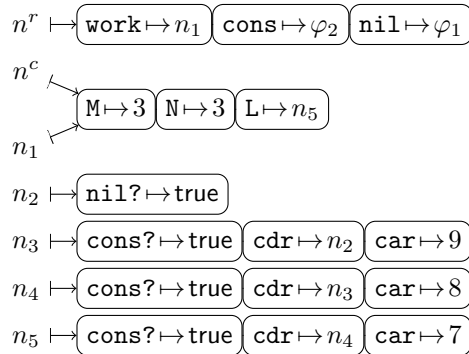
The inner loop call increments the counter by one and discards a single link. After the first recurrence, the current node appears as follows.



After the third recurrence, the procedure ends.



Finally, the caller's current node variable is restored and the current node variable M receives a copy of N. We lose the call's working node because no other references to it exists. Figure 1 shows the resulting graph-store. Here is the final heap.



3.2 Queues

Queues are a handy data structure for, among other things, distributed task scheduling. We can construct queues readily from pairs of lists [9] according to the banker's method. In this method, we call one list the *front* and the other the *rear*. New elements are *snoiced* onto the *rear*. Whenever the rear becomes longer than the front, we append the front to the reversed rear. Thus, we need list reverse and append operations.

The following implementation of the reverse operation simply copies the elements of list xs onto a new list ys in a tail recursive style. Because I use call to initiate the inner loop procedure, I do not need to give ys as an argument to loop. Despite the apparent lexical scoping of ys, pGrasp enforces no single variable scoping policy beyond the nature of nodes and paths. We might say that pGrasp encourages a policy of selective dynamic scope.




Figure 1: A list structure.

```

proc root reverse (xs) {
  new this ys nil ();
  proc this loop (xs) {
    when (not xs.nil?) {
      new this ys cons (xs.car,ys);
      call loop (xs.cdr) } };
  call loop (xs) }

```

We also need a list append operation. The next procedure first copies the reversed elements of *ys*, followed by the reversed elements of *xs*, onto a new list *zs*. An imperative style produces the most compact code for this example.

```

proc root append (xs,ys) {
  new this zs nil ();
  return this ys reverse (ys) ys;
  return this xs reverse (xs) ys;
  while (not ys.nil?) {
    new this zs cons (ys.car,zs);
    set this ys ys.cdr };
  while (not xs.nil?) {
    new this zs cons (xs.car,zs);
    set this xs xs.cdr } };

```

We are ready to implement the queue. The following constructor takes no arguments and constructs two lists with corresponding length counters. It also installs two helper procedures like OOP methods. The first, *snoc*, conses an element onto the rear of the queue and then calls the second, *cleanup*, to enforce the constraint on member list lengths.

```

proc root queue () {
  new this F nil ();
  new this R nil ();
  set this lenF 0;
  set this lenR 0;
  proc this snoc (x) {
    new this R cons (x,R);
    set this R (R + 1);
    call cleanup () } };
  proc this cleanup () {
    when (lenR > lenF) {
      return this R reverse (R) ys;
      return this F append (F,R) zs;
      new this R nil ();
      set this lenF (lenF + lenR);
      set this lenR 0 } }

```

Because the effects of *snoc* and *cleanup* are confined entirely to the queue instance they inhabit, the object-oriented style admits the simplest code. The same is not true for the remaining mutators because they return some result to the caller. The head procedure copies the element at the front of queue *q* into the variable *x* so the caller can fetch it with a return statement. If *q* is empty, *x* is null.

```

proc root head (q) {
  when (not q.F.nil?)
    set this x q.F.car }

```

The *tail* procedure duplicates the node referenced by *q* into variable *r*, then eliminates the front-most element of *r*. I use the *build* statement to direct *r.cleanup* into *r*. Equivalently, I could have built *q.cleanup* inside *r*.

```

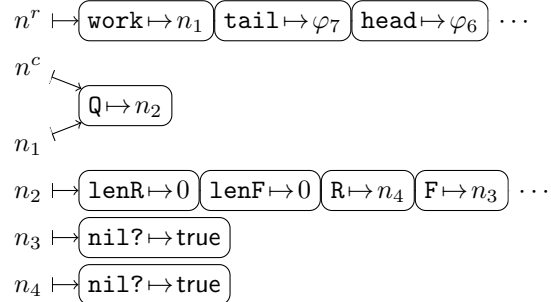
proc root tail (q) {
  if (q.F.nil?)
    set this r q;
  else {
    dup r this q;
    set r F F.cdr;
    set r lenF (r.lenF - 1);
    build r r.cleanup () } };

```

Now, suppose we start a new pGrasp session with lists and queues, and a fresh work space in the current node. First, construct an empty queue.

```
new this Q queue ();
```

The heap now has the following structure.



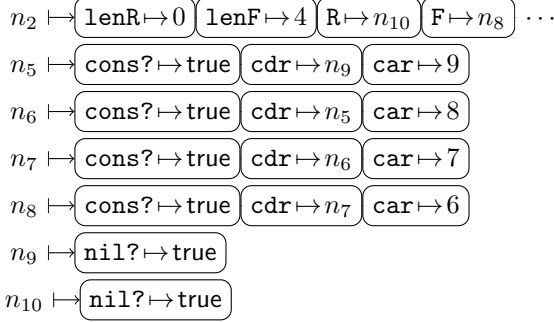
Enqueue four numbers.

```

build Q Q.snoc (6);
build Q Q.snoc (7);
build Q Q.snoc (8);
build Q Q.snoc (9);

```

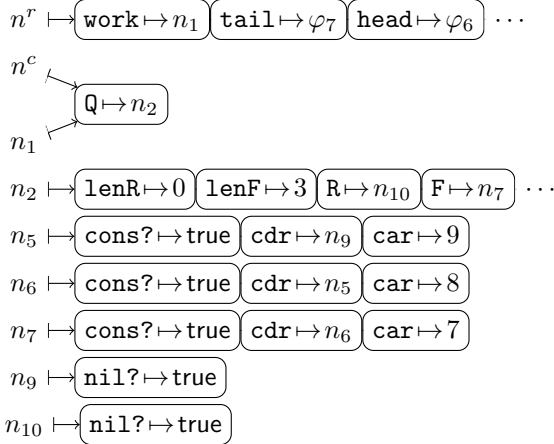

The first `snoc` conses 9 onto the rear of `Q`. Since `Q.F` is empty, `cleanup` moves the rear to the front. The second `snoc` conses 8 onto `Q.R` without calling `cleanup`. The third `snoc` behaves similar to the first.



Pop an element off the queue. Since `Q.R` is empty, no `cleanup` is necessary.

```
return this Q tail (Q) r;
```

The queue appears to shrink because the variable `Q` receives a copy of its “tail.” Figure 2 show the resulting graph-store. Here is the final heap.



The loose connectivity of this distributed task queue exemplifies a growing class of socially-motivated distributed software systems, where the operator of a system can contribute computing resources to others. With reasonable effort, we could extend this example with identity management and task submission quotas. Perhaps we could tie quotas meaningfully to the amount of resources contributed over time, to be bought and sold like a commodity, and enabling a digital equivalent to taxes.

3.3 Distributed Tasks

Since pGrasp does not support concurrency directly, I implement a distributed task scheduler in the following round-about manner. A dedicated pGrasp instance hosts a distributed task queue, along with a server program and a client API. The server program listens for incoming requests. The client API contains commands to request and submit labeled tasks and results.

The server’s task queue remains empty until another pGrasp instance submits some tasks. If an instance requests a task and the queue is not empty, the head task is removed from the queue and given to the requester along with a serial number. The requester must then run the task and submit the result back to the server. The server caches results until another instance requests them.

Suppose we have two descriptively named procedures, `listen` and `connect`, that take no arguments and produce a cursor named `c`. Then figure 3 contains an implementation of the distributed task queue.

The task submission command takes a list of nodes, where each node in the list contain a `label` variable and a run procedure that takes no arguments and stores its result in a variable named `result`. The results request command takes a label and stores any returned results in a variable named `results`. The task request command takes no arguments. It requests and runs a single task from the server, then returns the result back to the server along with the serial and label embedded in the task. The server program takes no arguments. It listens for a client action and reacts accordingly.

4. Formal Specification

4.1 Syntax

Figures 4-8 present the formal grammar. I partition the grammar into three major syntactic classes: statements, expressions, and values. A valid statement reduces to ε (skip) while possibly affecting the run-time state. A valid expression evaluates, without effect, to a *value*, i.e., a normal form. Denote statements, expressions, and values by meta-variables s , e , and v , respectively.

4.2 Semantics

Figures 9-12 present the operational semantics of pGrasp as inference rules in a small-step style. Denote by G a graph store, and by D a dispatch object. A judgment of the form $\langle G, e \rangle \rightsquigarrow e'$ specifies a step from expression e to expression e' in the context of graph-store G . A judgment of the form $\langle G, D, s \rangle \rightsquigarrow \langle G', D', s' \rangle$ specifies a step from statement s to statement s' in the context of graph-store G and dispatch object D . The resulting state consists of graph-store G' and dispatch object D' .

Denote by H , Σ , and n a heap and a node, and a node reference, respectively. Denote by \emptyset and Q a fresh empty node and a fresh empty queue, respectively. Denote by n^r and n^c a root reference and a current node reference, respectively. I occasionally decompose a graph-store into a tuple of its components: $G = (H, n^r, n^c)$.

Write Δ a dictionary such as a node or dispatch object. Denote by $\Delta[\alpha \mapsto \beta]$ the dictionary Δ extended with a mapping from α to β , where α is some reference object or variable name, depending on the nature of Δ , and β is some




Figure 2: A queue structure.

```

proc root sendtasks (label,tasks) {
  return this c connect () c;
  print c "send tasks";
  print c label;
  while (not tasks.nil?) {
    print c tasks.car;
    set this tasks tasks.cdr };
  print c "done";
  close c };

proc root getresults (label) {
  return this c connect () c;
  print c "get results";
  print c label;
  read c this results;
  close c };

proc root gettask () {
  return this c connect () c;
  print c "get task";
  read c this task;
  close c;
  when (not (task == null))
    build task task.run ();
  return this c connect () c;
  print c "send result";
  print c task;
  close c };

proc root server () {
  node this R;
  new this Q queue ();
  set this S = 1;
  while (true) {
    return this c listen () c;
    read c cmd;

    if (cmd == "send tasks") {
      read c task;
      while (not (task == "done")) {
        set task serial S;
        set this S (S + 1);
        build Q Q.snoc (task);
        read c task } }

    else if (cmd == "get task") {
      return this task head (Q) x;
      return this Q tail (Q) r;
      print c task }

    else if (cmd == "send result") {
      read c label;
      read c task;
      when (not R.label)
        new R label nil ();
        new R label cons (task, R.label) }

    else if (cmd == "get results") {
      read c label;
      print c R.label;
      unset R label } } }
  
```

Figure 3: A distributed task queue.

object suitable for storage in Δ . The equation $\Delta = \Delta'[\alpha \mapsto \beta]$ denotes a dictionary Δ' of all the mappings of Δ except the one from α , if it exists. Denote by $\Delta(\alpha)$ a dictionary look up:

$$\Delta(\alpha) = \beta \iff \Delta = \Delta'[\alpha \mapsto \beta]$$

If Δ contains no mapping from α , then $\Delta(\alpha) = \text{null}$. Denote by $\Delta(\alpha_1)(\alpha_2)$ a two step look up:

$$\Delta(\alpha_1)(\alpha_2) = \beta \iff \Delta(\alpha_1) = \Delta' \wedge \Delta'(\alpha_2) = \beta$$

For convenience, define the following single-variable update relation:

$$\Delta[\alpha_1:\alpha_2 \mapsto \beta] \equiv \Delta[\alpha_1 \mapsto \Delta(\alpha_1)[\alpha_2 \mapsto \beta]]$$

Denote by $\Delta_1 \cup \Delta_2$ the dictionary that contains all mappings of Δ_1 and Δ_2 , where the mappings of Δ_2 take precedence. Assuming Δ contains a mapping from α to some queue, denote by $\Delta[\alpha \Leftarrow \beta]$ the enqueue operation of β onto the queue referenced by α , and by $\Delta[\alpha \Rightarrow \beta]$ the dequeue operation of β from the queue referenced by α .

4.2.1 Helpers

Denote by $\text{args}(e, v)$ a fresh node in which the parameters of e are assigned to the corresponding arguments of v . Denote by $\text{copy}(n)$ a fresh node reference to a shallow copy of the node referenced by n . Denote by $\text{request}(n)$ a fresh cursor to the resource described by the node referenced by n , or null if the resource is unavailable. Denote by $\text{flush}(c)$ the queue referenced by cursor c , emptied by unspecified means. Denote by $\text{encode}(v)$ and $\text{decode}(v)$ the serialized and unserialized representation, respectively, of value v .

5. Implementation

I implemented the prototype in Racket [3, Version 5.3.6], a Scheme variant. The implementation is, however, behind the current specification, so some of the programs in this document have not been tested live.

In the implementation, I represent the skip statement by the empty string. For instance, in the statement,

```
while true;
```

the body of the loop is a single skip statement. For heap garbage collection, I track node linkage by simple reference counting. For term encoding and decoding, I use the built-in Racket term printer and reader. I transmit most primitive values as-is, omitting cursors, but effective node reference encoding is still an open question. For now, I encode referenced nodes recursively.

6. Discussion

The syntax of pGrasp is closest to C, but pGrasp is closer in spirit to assembly language—neither enforces a particularly complex call stack discipline and both require careful

though to avoid unintentional clobbering of data. Additionally, pGrasp has some elements of Java, e.g., references and garbage collection, along with the print/read constructs of Scheme. Because of the apparently low level at which the most basic pGrasp constructs operate, I believe a practical pGrasp platform with excellent performance, or at least reasonably good performance, is within reach. I can not yet predict the performance of a rewrite-based GSP language, although I think we could reasonably expect performance comparable to a purely functional language due to the maturity of the underlying theories and the performance of existing rewrite-based languages.

At a higher level of detail, pGrasp requires many constructs of convenience in order to produce reasonably compact code—four primitive procedure call constructs seems like too many. I also rely heavily on unscalable naming and calling conventions. I could alleviate these symptoms with a proper abstraction mechanism, but a basic macro facility could have similar effect and be more useful. Neither solution is ideal because neither is simple and none of these issues exist in a proper rewrite-based approach. Though staged evaluation might provide a scalable solution for returning values from procedure calls, the outcome betrays the overall vision for GSP and so I can not justify a possibly messy exploration into imperative staging.

The potential for explicit, high level call stack management in pGrasp is interesting in its own right, but a GSP setting offers benefits that are difficult or impossible to achieve in an imperative language. For example, fully pluggable type systems are all but trivial in a GSP language.

As a final thought, in either pGrasp or a pure GSP language, we could hook native programs directly into the encoding and decoding mechanisms with a simple callback mechanism, which would simplify integration with virtually all other technologies that use messaging, in the loose sense of messaging defined in the introduction.

7. Conclusion

I started the pGrasp project as an exercise in comparing GSP concepts to existing programming language technologies, to discover concrete examples from which to start a discussion about GSP and the issues it addresses. Designing a language like pGrasp, however, is like driving a screw with a sledge hammer—I get my point across, but I destroy the foundation in the process.

There is a clear discrepancy between the short-term goals of this project and the broader vision of the graph-store programming paradigm. With that in mind, I have identified at least three possible paths forward.

The GSP paradigm requires a fundamentally non-modular approach to software construction. Considering the pervasive importance of modularity in modern software engineering,

I do not expect rapid adoption of the paradigm until I put significant work into a production platform and supporting tools. While implementing the entire vision is clearly too large a task for a single Masters thesis, a coherent language based on graph rewriting—a staged, pattern-based design—is certainly within reach.

On the other hand, GSP principles may be applied, to some extent, in any language with dictionaries, including C, Python, and Perl. For each language, I could implement the pGrasp constructs in a cross-platform message processing and passing library, though performance of such a library in any language remains unclear. From a usability perspective, this approach is most attractive for languages with strong meta-programming or other DSL construction features, such as Scheme, Haskell, and Ruby.

The current pGrasp design alludes to a similar but more intuitive and performant revision. With some work, I could polish the design and produce a useful and reasonably efficient platform. The prospect of a practical pGrasp implementation has already drawn interest from members of academia and industry.

References

- [1] T. Brus, M. C. van Eekelen, M. Van Leer, and M. J. Plasmeijer. Cleana language for functional graph rewriting. In *Functional Programming Languages and Computer Architecture*, page 364384, 1987.
- [2] T. L. CASAVANT and J. G. KUHL. A taxonomy of scheduling in general-purpose distributed computing systems. *IEEE TRANSACTIONS ON SOFTWARE ENGINEERING*, 14(2):141, 1988.
- [3] M. Flatt and PLT. Reference: Racket. Technical Report PLT-TR-2010-1, PLT Inc., 2010. <http://racket-lang.org/tr1/>.
- [4] A. Goldberg and A. Kay. *Smalltalk-72: Instruction Manual*. Xerox Corporation, 1976.
- [5] A. Graf. *The Pure programming language*. See, 2009.
- [6] T. Isakowitz, M. Bieber, and F. Vitali. Web information systems. *Communications of the ACM*, 41(7):7880, 1998.
- [7] C. B. Jay. The pattern calculus. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 26(6):911937, 2004.
- [8] J. Meseguer. Conditional rewriting logic as a unified model of concurrency. *Theoretical computer science*, 96(1):73155, 1992.
- [9] C. Okasaki. *Purely functional data structures*. PhD thesis, Citeseer, 1996.
- [10] D. Plump. Term graph rewriting. *Handbook of graph grammars and computing by graph transformation*, 2:361, 1999.
- [11] G. Rozenberg. *Handbook of graph grammars and computing by graph transformation, vol 1: Foundations*. World Scientific, 1997.
- [12] E. Scott and A. Johnstone. GLL parsing. *Electronic Notes in Theoretical Computer Science*, 253(7):177–189, Sept. 2010.

$s ::= \varepsilon$	skip
$s; s$	sequence
$\text{node } e x$	create fresh node
$\text{dup } e e x$	duplicate node
$\text{set } e x e$	set variable
$\text{unset } e x$	drop variable
$\text{reset } e$	empty existing node
$\text{in } e$	change work node
$e ::= x$	variable
$e.e$	path expression
root	root node
this	current node
$v ::= n$	node reference
null	null reference

Figure 4: Core syntax.

$s ::= \dots$	
$\text{proc } e x e s$	define procedure
$\text{call } e e$	call procedure
$\text{build } e e e$	call-in-node
$\text{new } e x e e$	call-in-new-node
$\text{return } e x e e x$	call-and-save
$e ::= \dots$	
$()$	unit expression
e, e	tuple expression
$v ::= \dots$	
$()$	unit value
v, v	tuple value
$\varphi(e, s)$	procedure object

Figure 5: Procedure syntax.

ISSN 15710661. .

- [13] W. Taha. *Multi-stage programming: Its theory and applications*. PhD thesis, Citeseer, 1999.

$\frac{\text{E-PATH1} \quad \langle G, e_1 \rangle \rightsquigarrow e'_1}{\langle G, e_1.e_2 \rangle \rightsquigarrow e'_1.e_2}$	$\frac{\text{E-PATH21} \quad \langle \langle H, n^r, n \rangle, e \rangle \rightsquigarrow e'}{\langle G, n.e \rangle \rightsquigarrow n.e'}$	$\frac{\text{E-PATH22}}{\langle G, n_1.n_2.e \rangle \rightsquigarrow n_2.e}$	$\frac{\text{E-PATH}}{\langle G, n.v \rangle \rightsquigarrow v}$	$\frac{\text{E-VARIABLE}}{\langle G, x \rangle \rightsquigarrow G(n^c)(x)}$
$\frac{\text{E-ROOT}}{\langle G, \text{root} \rangle \rightsquigarrow n^r}$	$\frac{\text{E-THIS}}{\langle G, \text{this} \rangle \rightsquigarrow n^c}$	$\frac{\text{S-SEQUENCE1} \quad \langle G, D, s_1 \rangle \rightsquigarrow \langle G, D, s'_1 \rangle}{\langle G, D, s_1; s_2 \rangle \rightsquigarrow \langle G, D, s'_1; s_2 \rangle}$	$\frac{\text{S-SEQUENCE2}}{\langle G, D, \varepsilon; s \rangle \rightsquigarrow \langle G, D, s \rangle}$	
$\frac{\text{S-NODE1} \quad \langle G, e \rangle \rightsquigarrow e'}{\langle G, D, \text{node } e \ x \rangle \rightsquigarrow \langle G, D, \text{node } e' \ x \rangle}$	$\frac{\text{S-NODE} \quad G' = G[n:x \mapsto \emptyset]}{\langle G, D, \text{node } n \ x \rangle \rightsquigarrow \langle G', D, \varepsilon \rangle}$	$\frac{\text{S-DUP1} \quad \langle G, e_1 \rangle \rightsquigarrow e'_1}{\langle G, D, \text{dup } e_1 \ e_2 \ x \rangle \rightsquigarrow \langle G, D, \text{dup } e'_1 \ e_2 \ x \rangle}$		
$\frac{\text{S-DUP2} \quad \langle G, e \rangle \rightsquigarrow e'}{\langle G, D, \text{dup } n \ e \ x \rangle \rightsquigarrow \langle G, D, \text{dup } n \ e' \ x \rangle}$		$\frac{\text{S-DUP} \quad G' = G[n_2:x \mapsto \text{copy}(n_1)]}{\langle G, D, \text{dup } n_1 \ n_2 \ x \rangle \rightsquigarrow \langle G', D, \varepsilon \rangle}$		
$\frac{\text{S-SET1} \quad \langle G, e_1 \rangle \rightsquigarrow e'_1}{\langle G, D, \text{set } e_1 \ x \ e_2 \rangle \rightsquigarrow \langle G, D, \text{set } e'_1 \ x \ e_2 \rangle}$	$\frac{\text{S-SET3} \quad \langle G, e \rangle \rightsquigarrow e'}{\langle G, D, \text{set } n \ x \ e \rangle \rightsquigarrow \langle G, D, \text{set } n \ x \ e' \rangle}$	$\frac{\text{S-SET} \quad G' = G[n:x \mapsto v]}{\langle G, D, \text{set } n \ x \ v \rangle \rightsquigarrow \langle G', D, \varepsilon \rangle}$		
$\frac{\text{S-UNSET1} \quad \langle G, e \rangle \rightsquigarrow e'}{\langle G, D, \text{unset } e \ x \rangle \rightsquigarrow \langle G, D, \text{unset } e' \ x \rangle}$	$\frac{\text{S-UNSET} \quad G(n) = \Sigma[x \mapsto v]}{\langle G, D, \text{unset } n \ x \rangle \rightsquigarrow \langle G[n \mapsto \Sigma], \varepsilon \rangle}$	$\frac{\text{S-RESET1} \quad \langle G, e \rangle \rightsquigarrow e'}{\langle G, D, \text{reset } e \rangle \rightsquigarrow \langle G, D, \text{reset } e' \rangle}$		
$\frac{\text{S-RESET} \quad G' = G[n \mapsto \emptyset]}{\langle G, D, \text{reset } n \rangle \rightsquigarrow \langle G', D, \varepsilon \rangle}$	$\frac{\text{S-IN1} \quad \langle G, e \rangle \rightsquigarrow e'}{\langle G, D, \text{in } e \rangle \rightsquigarrow \langle G, D, \text{in } e' \rangle}$	$\frac{\text{S-IN}}{\langle G, D, \text{in } n \rangle \rightsquigarrow \langle \langle H, n^r, n \rangle, D, \varepsilon \rangle}$		

Figure 9: Core semantics.

$s ::= \dots$ open $e \ e \ x$ close e send $e \ e$ rcv $e \ e \ x$ print $e \ e$ read $e \ e \ x$	create cursor destroy cursor raw output raw input structured output structured input	$s ::= \dots$ when $e \ s$ if $e \ s$ else s while $e \ s$	optional statement conditional statement iteration
$e ::= \dots$ stdio stderr	standard input/output standard error	$e ::= \dots$ if $e \ e$ else e e and e e or e not e $e == e$ true false	conditional expression conjunction disjunction negation equality test truth untruth
$v ::= \dots$ c	cursor value	$v ::= \dots$ true false	truth value untruth value

Figure 6: I/O syntax.

Figure 7: Logic syntax.

$$\begin{array}{c}
\text{E-TUPLE1} \\
\frac{\langle G, e_1 \rangle \rightsquigarrow e'_1}{\langle G, e_1, e_2 \rangle \rightsquigarrow e'_1, e_2} \\
\\
\text{E-TUPLE2} \\
\frac{\langle G, e \rangle \rightsquigarrow e'}{\langle G, v, e \rangle \rightsquigarrow v, e'} \\
\\
\text{S-PROC1} \\
\frac{\langle G, e_1 \rangle \rightsquigarrow e'_1}{\langle G, D, \text{proc } e_1 \ x \ e_2 \ s \rangle \rightsquigarrow \langle G, D, \text{proc } e'_1 \ x \ e_2 \ s \rangle} \\
\\
\text{S-PROC} \\
\frac{G' = G[n:x \mapsto \varphi(e, s)]}{\langle G, D, \text{proc } n \ x \ e \ s \rangle \rightsquigarrow \langle G', D, \varepsilon \rangle} \\
\\
\text{S-CALL1} \\
\frac{\langle G, e_1 \rangle \rightsquigarrow e'_1}{\langle G, D, \text{call } e_1 \ e_2 \rangle \rightsquigarrow \langle G, D, \text{call } e'_1 \ e_2 \rangle} \\
\\
\text{S-CALL2} \\
\frac{\langle G, e \rangle \rightsquigarrow e'}{\langle G, D, \text{call } \varphi \ e \rangle \rightsquigarrow \langle G, D, \text{call } \varphi \ e' \rangle} \\
\\
\text{S-CALL} \\
\frac{G' = G[n^c \mapsto G(n^c) \cup \text{args}(e, v)]}{\langle G, D, \text{call } \varphi(e, s) \ v \rangle \rightsquigarrow \langle G', D, s \rangle} \\
\\
\text{S-BUILD1} \\
\frac{\langle G, e_1 \rangle \rightsquigarrow e'_1}{\langle G, D, \text{build } e_1 \ e_2 \ e_3 \rangle \rightsquigarrow \langle G, D, \text{build } e'_1 \ e_2 \ e_3 \rangle} \\
\\
\text{S-BUILD2} \\
\frac{\langle G, e_2 \rangle \rightsquigarrow e'_2}{\langle G, D, \text{build } n \ e_2 \ e_3 \rangle \rightsquigarrow \langle G, D, \text{build } n \ e'_2 \ e_3 \rangle} \\
\\
\text{S-BUILD3} \\
\frac{\langle G, e_3 \rangle \rightsquigarrow e'_3}{\langle G, D, \text{build } n \ \varphi \ e_3 \rangle \rightsquigarrow \langle G, D, \text{build } n \ \varphi \ e'_3 \rangle} \\
\\
\text{S-BUILD} \\
\frac{s = (\text{in } n; \text{call } \varphi \ v; \text{in } n^c)}{\langle G, D, \text{build } n \ \varphi \ v \rangle \rightsquigarrow \langle G, D, s \rangle} \\
\\
\text{S-NEW} \\
\frac{s = (\text{node } e_1 \ x_1; \text{build } e_1.x_1 \ e_2 \ e_3)}{\langle G, D, \text{new } e_1 \ x_1 \ e_2 \ e_3 \rangle \rightsquigarrow \langle G, D, s \rangle} \\
\\
\text{E-RETURN} \\
\frac{s = (\text{build } n \ e_2 \ e_3; \text{set } e_1 \ x_1 \ n.x_2) \quad n \text{ fresh}}{\langle G, D, \text{return } e_1 \ x_1 \ e_2 \ e_3 \ x_2 \rangle \rightsquigarrow \langle G, D, s \rangle}
\end{array}$$

Figure 10: Procedure semantics.

$e ::= e + e$	addition
$e - e$	subtraction
$e * e$	multiplication
e / e	division
\mathbb{R}	number
$v ::= \dots$	
\mathbb{R}	numeric value

Figure 8: Arithmetic syntax.

$$\begin{array}{c}
\text{E-STDIO} \quad \frac{}{\langle G, \text{stdio} \rangle \rightsquigarrow c^{io}} \quad \text{E-STDERR} \quad \frac{}{\langle G, \text{stderr} \rangle \rightsquigarrow c^{err}} \quad \text{S-OPEN1} \quad \frac{\langle G, e_1 \rangle \rightsquigarrow e'_1}{\langle G, D, \text{open } e_1 e_2 x \rangle \rightsquigarrow \langle G, D, \text{open } e'_1 e_2 x \rangle} \\
\text{S-OPEN2} \quad \frac{\langle G, e \rangle \rightsquigarrow e'}{\langle G, D, \text{open } n e x \rangle \rightsquigarrow \langle G, D, \text{open } n e' x \rangle} \quad \text{S-OPEN} \quad \frac{G' = G[n_2:x \mapsto c] \quad D'[c \mapsto Q] \quad c = \text{request}(n_1)}{\langle G, D, \text{open } n_1 n_2 x \rangle \rightsquigarrow \langle G', D', \varepsilon \rangle} \\
\text{S-CLOSE1} \quad \frac{\langle G, e \rangle \rightsquigarrow e'}{\langle G, D, \text{close } e \rangle \rightsquigarrow \langle G, D, \text{close } e' \rangle} \quad \text{S-CLOSE} \quad \frac{D = D'[c \mapsto \text{flush}(c)]}{\langle G, D, \text{close } c \rangle \rightsquigarrow \langle G, D', \varepsilon \rangle} \quad \text{S-SEND1} \quad \frac{\langle G, e_1 \rangle \rightsquigarrow e'_1}{\langle G, D, \text{send } e_1 e_2 \rangle \rightsquigarrow \langle G, D, \text{send } e'_1 e_2 \rangle} \\
\text{S-SEND2} \quad \frac{\langle G, e \rangle \rightsquigarrow e'}{\langle G, D, \text{send } c e \rangle \rightsquigarrow \langle G, D, \text{send } c e' \rangle} \quad \text{S-SEND} \quad \frac{D' = D[c \Leftarrow v]}{\langle G, D, \text{send } c v \rangle \rightsquigarrow \langle G, D', \varepsilon \rangle} \quad \text{S-RCV1} \quad \frac{\langle G, e_1 \rangle \rightsquigarrow e'_1}{\langle G, D, \text{recv } e_1 e_2 x \rangle \rightsquigarrow \langle G, D, \text{recv } e'_1 e_2 x \rangle} \\
\text{S-RCV2} \quad \frac{\langle G, e \rangle \rightsquigarrow e'}{\langle G, D, \text{recv } c e x \rangle \rightsquigarrow \langle G, D, \text{recv } c e' x \rangle} \quad \text{S-RCV} \quad \frac{D' = D[c \Rightarrow v] \quad G' = G[n:x \mapsto v]}{\langle G, D, \text{recv } c n x \rangle \rightsquigarrow \langle G', D', \varepsilon \rangle} \\
\text{S-PRINT1} \quad \frac{\langle G, e_1 \rangle \rightsquigarrow e'_1}{\langle G, D, \text{print } e_1 e_2 \rangle \rightsquigarrow \langle G, D, \text{print } e'_1 e_2 \rangle} \quad \text{S-PRINT2} \quad \frac{\langle G, e \rangle \rightsquigarrow e'}{\langle G, D, \text{print } c e \rangle \rightsquigarrow \langle G, D, \text{print } c e' \rangle} \quad \text{S-PRINT} \quad \frac{D' = D[c \Leftarrow \text{encode}(v)]}{\langle G, D, \text{print } c v \rangle \rightsquigarrow \langle G, D', \varepsilon \rangle} \\
\text{S-READ1} \quad \frac{\langle G, e_1 \rangle \rightsquigarrow e'_1}{\langle G, D, \text{read } e_1 e_2 x \rangle \rightsquigarrow \langle G, D, \text{read } e'_1 e_2 x \rangle} \quad \text{S-READ2} \quad \frac{\langle G, e \rangle \rightsquigarrow e'}{\langle G, D, \text{read } c e x \rangle \rightsquigarrow \langle G, D, \text{read } c e' x \rangle} \\
\text{S-READ} \quad \frac{D' = D[c \Rightarrow v] \quad G' = G[n:x \mapsto \text{decode}(v)]}{\langle G, D, \text{read } c n x \rangle \rightsquigarrow \langle G', D', \varepsilon \rangle}
\end{array}$$

Figure 11: I/O semantics.

$$\begin{array}{c}
\text{E-IF1} \quad \frac{\langle G, e_1 \rangle \rightsquigarrow e'_1}{\langle G, \text{if } e_1 \ e_2 \ \text{else } e_3 \rangle \rightsquigarrow \text{if } e'_1 \ e_2 \ \text{else } e_3} \qquad \text{E-IF2} \quad \frac{v \notin \{\text{false}, \text{null}\}}{\langle G, \text{if } v \ e_2 \ \text{else } e_3 \rangle \rightsquigarrow e_2} \qquad \text{E-IF3} \quad \frac{v \in \{\text{false}, \text{null}\}}{\langle G, \text{if } v \ e_2 \ \text{else } e_3 \rangle \rightsquigarrow e_3} \\
\\
\text{E-AND1} \quad \frac{\langle G, e_1 \rangle \rightsquigarrow e'_1}{\langle G, e_1 \ \text{and } e_2 \rangle \rightsquigarrow e'_1 \ \text{and } e'_2} \qquad \text{E-AND2} \quad \frac{v \notin \{\text{false}, \text{null}\}}{\langle G, v \ \text{and } e_2 \rangle \rightsquigarrow e_2} \qquad \text{E-AND3} \quad \frac{v \in \{\text{false}, \text{null}\}}{\langle G, v \ \text{and } e_2 \rangle \rightsquigarrow \text{false}} \qquad \text{E-OR1} \quad \frac{\langle G, e_1 \rangle \rightsquigarrow e'_1}{\langle G, e_1 \ \text{or } e_2 \rangle \rightsquigarrow e'_1 \ \text{or } e_2} \\
\\
\text{E-OR2} \quad \frac{v \in \{\text{false}, \text{null}\}}{\langle G, v \ \text{or } e_2 \rangle \rightsquigarrow e_2} \qquad \text{E-OR3} \quad \frac{v \notin \{\text{false}, \text{null}\}}{\langle G, v \ \text{or } e_2 \rangle \rightsquigarrow v} \qquad \text{E-NOT} \quad \frac{\langle G, e \rangle \rightsquigarrow e'}{\langle G, \text{not } e \rangle \rightsquigarrow \text{not } e'} \qquad \text{E-NOT1} \quad \frac{v \in \{\text{false}, \text{null}\}}{\langle G, \text{not } v \rangle \rightsquigarrow \text{true}} \qquad \text{E-NOT2} \quad \frac{v \notin \{\text{false}, \text{null}\}}{\langle G, \text{not } v \rangle \rightsquigarrow \text{false}} \\
\\
\text{E-EQUALS1} \quad \frac{\langle G, e_1 \rangle \rightsquigarrow e'_1}{\langle G, e_1 == e_2 \rangle \rightsquigarrow e'_1 == e_2} \qquad \text{E-EQUALS2} \quad \frac{\langle G, e_2 \rangle \rightsquigarrow e'_2}{\langle G, v_1 == e_2 \rangle \rightsquigarrow v_1 == e'_2} \qquad \text{E-EQUALS3} \quad \frac{v_1 = v_2}{\langle G, v_1 == v_2 \rangle \rightsquigarrow \text{true}} \qquad \text{E-EQUALS4} \quad \frac{v_1 \neq v_2}{\langle G, v_1 == v_2 \rangle \rightsquigarrow \text{false}} \\
\\
\text{S-WHEN} \quad \frac{\langle G, e \rangle \rightsquigarrow e'}{\langle G, D, \text{when } e \ s \rangle \rightsquigarrow \langle G, D, \text{when } e' \ s \rangle} \qquad \text{S-WHEN1} \quad \frac{v \in \{\text{false}, \text{null}\}}{\langle G, D, \text{when } v \ s \rangle \rightsquigarrow \langle G, D, \varepsilon \rangle} \qquad \text{S-WHEN2} \quad \frac{v \notin \{\text{false}, \text{null}\}}{\langle G, D, \text{when } v \ s \rangle \rightsquigarrow \langle G, D, s \rangle} \\
\\
\text{S-IF} \quad \frac{\langle G, e_1 \rangle \rightsquigarrow e'_1}{\langle G, D, \text{if } e \ s_1 \ \text{else } s_2 \rangle \rightsquigarrow \langle G, D, \text{if } e' \ s_1 \ \text{else } s_2 \rangle} \qquad \text{S-IF1} \quad \frac{v \notin \{\text{false}, \text{null}\}}{\langle G, D, \text{if } v \ s_1 \ \text{else } s_2 \rangle \rightsquigarrow \langle G, D, s_1 \rangle} \\
\\
\text{S-IF2} \quad \frac{v \in \{\text{false}, \text{null}\}}{\langle G, D, \text{if } v \ s_1 \ \text{else } s_2 \rangle \rightsquigarrow \langle G, D, s_2 \rangle} \qquad \text{S-WHILE} \quad \frac{}{\langle G, D, \text{while } e \ s \rangle \rightsquigarrow \langle G, D, \text{when } e \ (s; \text{while } e \ s) \rangle}
\end{array}$$

Figure 12: Logic semantics.