$\text{Pl}\omega\text{NK}$: Functional Probabilistic NetKAT

ALEXANDER VANDENBROUCKE, KU Leuven, Belgium

TOM SCHRIJVERS, KU Leuven, Belgium

This work presents $\text{Pl}\omega\text{NK}$, a functional probabilistic network programming language that extends Probabilistic NetKAT (PNK). Like PNK, it enables probabilistic modelling of network behaviour, by providing probabilistic choice and infinite iteration (to simulate looping network packets). Yet, unlike PNK, it also offers abstraction and higher-order functions to make programming much more convenient.

The formalisation of $\text{Pl}\omega\text{NK}$ is challenging for two reasons: Firstly, network programming induces multiple side effects (in particular, parallelism and probabilistic choice) which need to be carefully controlled in a functional setting. Our system uses an explicit syntax for thunks and sequencing which makes the interplay of these effects explicit. Secondly, measure theory, the standard domain for formalisations of (continuous) probabilistic languages, does not admit higher-order functions. We address this by leveraging ω -Quasi Borel Spaces (ωQBSes), a recent advancement in the domain theory of probabilistic programming languages.

We believe that our work is not only useful for bringing abstraction to PNK, but that—as part of our contribution—we have developed the meta-theory for a probabilistic language that combines advanced features like higher-order functions, iteration and parallelism, which may inform similar meta-theoretic efforts.

CCS Concepts: • **Networks**; • **Software and its engineering** → **Domain specific languages**; • **Mathematics of computing** → *Probability and statistics*;

Additional Key Words and Phrases: Probabilistic Programming, Network Modelling, Quasi-Borel Spaces, $\omega\text{-QBS}$, NetKAT

ACM Reference Format:

Alexander Vandenbroucke and Tom Schrijvers. 2020. $\text{Pl}\omega\text{NK}$: Functional Probabilistic NetKAT. *Proc. ACM Program. Lang.* 4, POPL, Article 39 (January 2020), 27 pages. <https://doi.org/10.1145/3371107>

1 INTRODUCTION

Probabilistic programming languages simplify the creation of probabilistic models. They separate the model from the algorithm that infers probabilities for it (e.g., Church [Goodman et al. 2012], Anglican [Wood et al. 2014], Gen [Cusumano-Towner et al. 2019], ProbLog [Fierens et al. 2015]). Instead of writing a custom procedure tailored to a particular model, the same generic algorithm is used for all programs written in the programming language. Thus, the algorithm can be re-used for many programs, lessening the implementation effort and maintenance burden of the model.

In this work we develop a probabilistic programming language, called $\text{Pl}\omega\text{NK}$.¹ $\text{Pl}\omega\text{NK}$ combines diverse features such as higher-order functions, probabilistic choice and parallelism. It is a domain specific language for probabilistically modelling computer networks. The main purpose of $\text{Pl}\omega\text{NK}$

¹Pronounced as “plonk”.

Authors’ addresses: Alexander Vandenbroucke, KU Leuven, Celestijnenlaan 200 A, 3001, Leuven, Belgium, alexander.vandenbroucke@kuleuven.be; Tom Schrijvers, KU Leuven, Celestijnenlaan 200 A, 3001, Leuven, Belgium, tom.schrijvers@kuleuven.be.

This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2020 Copyright held by the owner/author(s).

2475-1421/2020/1-ART39

<https://doi.org/10.1145/3371107>

is to model computer networks and network protocols at an abstract level, and verify a wide variety of properties of such models, e.g., latency, fault-tolerance, or the absence of routing loops. The motivation is the same as for formal verification of computer programs or the mechanical checking of proofs. Namely, during the design of complex networks or protocols, even the best designers are bound to make some mistakes or errors. A computer-checked specification can detect such deficiencies before they are deployed [Anderson et al. 2014; Foster et al. 2016].

$\text{Pl}\omega\text{NK}$ extends Probabilistic NetKAT (PNK) [Foster et al. 2016]. Here is a small PNK program:

$$\underbrace{(sw = 1; drop \oplus_{0.1} sw \leftarrow 2 \& sw = 2; drop \oplus_{0.1} sw \leftarrow 1)}_{\text{node 1}} \mid \underbrace{)}_{\text{node 2}} \quad \left| \quad \begin{array}{c} \text{1} \xleftrightarrow{10\%} \text{2} \end{array} \right.$$

The program on the left models the network on the right. This network consists of only two nodes, 1 and 2, with a bidirectional link between them. The link between the nodes is unreliable, e.g., it is a radio link with poor reception. This causes a 10% of the packets to be lost in transit.

At this point it is not important to understand the meaning of this program exactly. Instead, note that this example already features a lot of repetition: the sub-programs to the left and to the right of the $\&$ -operator—modelling the behaviour of node 1 or 2, respectively—essentially mirror each other. Unfortunately, PNK offers no facilities to take advantage of this insight. The key advantage of $\text{Pl}\omega\text{NK}$ over PNK is that we can exploit it by abstracting over the behaviour of both parts with a function:

$$\underbrace{forward = \lambda src. \lambda dst. (sw = src; drop \oplus_{0.1} sw \leftarrow dst)}_{\text{node 1}} \mid \underbrace{(forward\ 1\ 2 \& forward\ 2\ 1)}_{\text{node 2}} \quad \left| \quad \begin{array}{c} \text{1} \xleftrightarrow{10\%} \text{2} \end{array} \right.$$

The function *forward* captures the general forwarding behaviour of the nodes, independently of a particular node. While this change arguably does not make the program shorter, the improved readability and maintainability make writing and extending the program much more convenient. For instance, adding a third node is now much easier:

$$\underbrace{forward = \lambda src. \lambda dst. (sw = src; drop \oplus_{0.1} sw \leftarrow dst)}_{\text{node 1}} \mid \underbrace{(forward\ 1\ 2 \& forward\ 2\ 3 \& forward\ 3\ 1)}_{\text{node 2}} \mid \underbrace{)}_{\text{node 3}} \quad \left| \quad \begin{array}{c} \text{1} \xleftrightarrow{10\%} \text{2} \\ \text{1} \xleftrightarrow{10\%} \text{3} \\ \text{2} \xleftrightarrow{10\%} \text{3} \end{array} \right.$$

Explicit Syntax for Thunks and Sequencing. Three distinct side-effects are in evidence in the above examples: (1) state—as we explain later, $sw \leftarrow 1$ modifies packets; (2) parallelism—the subprograms for node 1, 2 and 3 are run in parallel with $\&$; and (3) probability—through the $\oplus_{0.1}$ -operator. We carefully chose the previous example such that no arguments to the function contained any side-effects. Indeed, all arguments were constants, either 1, 2 or 3. Although this is already quite useful, we want to be more flexible in our full language, and also apply functions to non-constant expressions. In this case, should the side-effects of this expression be executed before the application and only the resulting value passed to the function (*Call-By-Value*)? Or, should the expression remain unevaluated, allowing the function to decide when to evaluate it (*Call-By-Name*)? Both strategies have their merits and there is no clear winner.

Rather than fix any particular order, we choose to explicitly segregate expressions into *computations* (which have side-effects) and *values* (which do not), loosely inspired by Call-By-Push-Value [Levy 2001].

The sequencing of side-effects then becomes explicit: either the computation is evaluated, producing a value which is then passed to the function, or the computation is explicitly turned into a

value, by wrapping it in a *thunk*. While we use CBPV as an inspiration for the syntax and semantics, our language does not enjoy all theoretical properties of CBPV and thus does not model it exactly.

Semantics. Our language, $\text{Pl}\omega\text{NK}$, has higher-order functions, iteration and probabilistic choice. This significantly complicates the formalisation of the semantics of our language. After all, the standard approach to define denotational semantics for a continuous probabilistic language is based on measure theory. Yet, measure theory does not support general higher-order functions [Aumann et al. 1961], a central feature of our language. Thus, to support higher-order functions, we use (ω) -Quasi-Borel Spaces [Heunen et al. 2017; Vákár et al. 2019] as the domain of our denotational semantics. This is a recently developed alternative axiomatisation of probability theory, which admits higher-order functions. Moreover, it possesses the ωCPO structure required to model PNK style iteration (Kleene-star).

Another challenge we face is the interaction of higher-order functions, state, parallelism and probability. In our language, a computation can produce a function in a manner that is simultaneously probabilistic and non-deterministic (parallel), and also locally² modifying state. Moreover, we must avoid accidentally duplicating work of parallel branches, since parallel composition ($\&$) is not idempotent (i.e., a program p is *not* equivalent to $p \& p$). As a result, defining the correct semantics for function application and sequencing is complicated and highly non-trivial. Note that this challenge is present in any calculus that supports function application, probability, parallelism and state—in fact, it is somewhat easier in our setting, as side-effects cannot occur inside the arguments of a function application.

As we mentioned earlier, $\text{Pl}\omega\text{NK}$'s primary purpose is the *specification* of network models and the verification of properties of those models. It shares this purpose with PNK, which has several computational properties that make it well suited for this purpose: The denotational semantics of PNK can be approximated [Smolka et al. 2017b] computationally, through an iterative procedure. Moreover, at the cost of disallowing the *dup* operation, PNK has decidable program equivalence [Smolka et al. 2017a, 2019]. Thus, if we show that a small—hence, easy to prove correct—program is correct, we implicitly show that all equivalent larger programs are also correct.

We show that our language also possesses these properties, subject to some (minor) restrictions: the approximation procedure exists, if we forbid parallel choice between functions. Essentially, this restricts $\text{Pl}\omega\text{NK}$'s parallelism to the parallelism that is present in PNK. Also, we conjecture that disallowing the *dup* operation, as for PNK, results in decidable program equivalence for $\text{Pl}\omega\text{NK}$. The specific contributions of this work are:

- We define the probabilistic programming language $\text{Pl}\omega\text{NK}$, for modelling computer networks and protocols. It features higher-order functions, probabilistic choice and parallelism. $\text{Pl}\omega\text{NK}$ extends the earlier programming language PNK.
- $\text{Pl}\omega\text{NK}$ extends PNK with a simple type system. The type system is important, not only for rejecting invalid programs, but also to ensure that all programs *Strongly Normalise* [Pierce 2002]. On the one hand, strong normalisation indirectly makes our denotational semantics well-defined. On the other hand, we exploit this property for compiling $\text{Pl}\omega\text{NK}$ to PNK. Recall that $\text{Pl}\omega\text{NK}$ is a specification language, and the flexibility of general recursion and real arithmetic is not required.
- We define denotational semantics for $\text{Pl}\omega\text{NK}$. As $\text{Pl}\omega\text{NK}$ contains higher-order functions, iteration and probabilistic choice, we need to leverage recent advances in the domain theory for probabilistic programs by Vákár et al. [2019]. They define an alternative formalisation

²By local, we mean that the state is not shared between different parallel or probabilistic branches.

of probability theory that admits higher-order functions, and iterations, the ω -Quasi-Borel Spaces (ω QBSes).

We prove several well-definedness theorems about this denotational semantics. The (pen-and-paper) proofs for several of these theorems use logical relations whose definitions are interesting in their own right. We also prove that this semantics is a conservative extension of PNK's semantics as given by Smolka et al. [2017b]. The proofs can be found in the extended version of this paper.

- We develop a subclass of $\text{Pl}\omega\text{NK}$ programs which can be compiled into PNK. Through the type system, we restrict the parallelism in $\text{Pl}\omega\text{NK}$ to the parallelism present in PNK. This makes the values that are produced in parallel more predictable, allowing compilation to succeed. Moreover, PNK itself lies entirely within this class. We have mechanised the meta-theory of $\text{Pl}\omega\text{NK}$ and the compilation procedure with the aid of the Abella proof-assistant [Gacek 2008]. Theorems bearing a check mark (✓) have been mechanised. The proof scripts are available in the supplementary material (<https://bitbucket.org/AlexanderV/probnetkat-lambda/src/master/abella/>).

Compilation preserves the denotational semantics of $\text{Pl}\omega\text{NK}$. Since our semantics is a conservative extension, the compiled PNK program behaves identically to the original $\text{Pl}\omega\text{NK}$ source program. These theorems are not easily encoded in Abella, and thus have been proven the classical way, with pen and paper.

- We have implemented a prototype of $\text{Pl}\omega\text{NK}$ in Haskell (<https://bitbucket.org/AlexanderV/probnetkat-lambda/src/master/>). The prototype implements a small extension to $\text{Pl}\omega\text{NK}$ that performs type reconstruction. The implementation can be used to run the small examples that are presented in this article.

2 OVERVIEW

2.1 A Brief Introduction to PNK

The central notion of PNK are *packet histories*, i.e., ordered sequences of packets. We write $\pi :: h$ for a history consisting of its most recent packet π , followed by the earlier history h . The empty history is written as $\langle \rangle$. The set of all packet histories is denoted PH . PNK programs operate on sets of these histories.

Packets themselves are intended to be a simplified model of real-world binary network packets. As such, they consist of a number of header fields, which are assigned numeric values. Contrary to real-world packets, they do not contain a payload, because it is irrelevant for routing decisions. In our examples we commonly use the following headers: the switch the packet is currently at (sw) and the port the packet is currently at (pt).

PNK programs are constructed by composing a number of primitive operations. These primitives are predicates (e.g., $drop$ or $sw = 1$), assignments (e.g., $pt \leftarrow 2$), and duplication. Recall that PNK programs operate on sets of packets histories. *Predicates* filter this set, allowing only specific histories. For instance, tests such as $sw = 1$ only allow histories where the first packet's header sw is set to 1, whereas $drop$ denies all packets, producing an empty set. *Assignments* instead modify the histories in the set. For instance, the expression $pt \leftarrow 2$ sets the pt header to 2 for the first packet of every history. *Duplication* dup duplicates and prepends the first packet of every history in the set, i.e. for history $\pi :: h$, dup produces a history $\pi :: \pi :: h$.

Primitive operations can be composed in three ways:³ sequentially, parallelly or probabilistically. *Sequential composition* ($;$) executes both operations one after the other, the output of the first becoming the input of the second. For instance, $sw = 1 ; pt \leftarrow 2$ first filters out all histories where

³Actually, there is a distinction between composition for predicates and other operations, but it is not relevant here.

Main Expression

$$(p ; t)^* ; sw = 4$$

Fig. 1. PNK and Pl ω NK models of a small network consisting of 4 nodes.

the *sw* header of the first packet is not 1, and then modifies the *pt* header of the first packet of every history. *Parallel composition* (&) executes both operations independently and then takes the union of the resulting sets. For instance, $sw = 1 \ \& \ sw = 2$ allows only packet histories that have the first packet's header set to either 1 or 2. Finally, *probabilistic choice* ($\oplus_r$) chooses either the left side with probability r , or the right with probability $1 - r$. For instance $sw \leftarrow 1 \oplus_{0.5} sw \leftarrow 2$, probabilistically chooses to set the first packet's header to 1 or to 2. If we sample from this expression, we see a set where the first packets' headers are either all set to 1, or all set to 2. The probability of either event is exactly 50%.

2.2 Modelling in PNK

Let us consider how to model the network shown in the top-left corner of Figure 1 in PNK. The network consists of four nodes, with links from node 1 to nodes 2 and 3, and from nodes 2 and 3 to node 4. The objective is to send a network packet from node 1 to node 4, routed through either node 2 or node 3.

In order to accurately model the network, the program must model two independent aspects: the *network topology*, that is, the links connecting the nodes (and their behaviour), and the *routing programs* running on the nodes, receiving and forwarding incoming packets. The PNK program modelling the network is shown on the left of Figure 1.

Topology. The topology is captured by the term t on the left-hand side of Figure 1. It is a parallel composition of terms. It models, from top to bottom, the links from 1 to 2, from 1 to 3, from 2 to 4 and from 3 to 4. Each of the links consists of a number of statements, composed sequentially. For

each link, it is verified that a packet is actually at the origin of the link by a series of guards (e.g., $sw = 1 ; pt = 3$, for the link from 1 to 3). Then, the packet is modified, setting sw and pt to the next switch on the link (e.g. $sw \leftarrow 3 ; pt \leftarrow 1$), thus transmitting the packet across the link. For the link between 1 and 2, we model unreliability by making a probabilistic choice ($\oplus_{0.9}$), choosing normal transmission 90% of the time and dropping (with *drop*) the packet 10% of the time. Packets that do not match the switch and port are rejected before their headers are modified.

Routing. The routing program is captured by the term p , a parallel composition of the routing programs for each node. For each node, it is verified that the sw field matches the node, and a node-specific routing program is then run: at node 1, packets are forwarded to node 2 and at node 2 packets are forwarded to node 4. Node 3 is unused for now.

Main Expression. The main expression (at the bottom of Figure 1) combines topology and routing, and provides an exit predicate. The $(\cdot)^*$, Kleene star, means iteration. Thus, the program $p ; t$ is repeated until the exit predicate $sw = 4$ is satisfied. This predicate checks that the packet has arrived at node 4, its destination.

Having a PNK model of our network allows us to estimate the probability of certain queries, such as the probability that a packet reaches its destination (90% in this case) or measure the expected congestion. This is by no means an exhaustive list. A slightly modified program permits us to estimate the latency (i.e., the average length of a path), or by restricting the language, program equivalence becomes decidable, creating an easy way to verify correctness [Smolka et al. 2017a, 2019]. For additional examples and details, the interested reader should consult the earlier work of Foster et al. [2016].

2.3 Extending PNK with functions

Even the small example from the previous section is quite tedious and repetitive to write. The root cause of this issue is the complete lack of abstraction facilities in PNK, since it is well-known that abstraction improves modularity and enables code re-use.

Arguably one of the most basic abstraction facilities available in programming languages is the venerable λ -abstraction. In $\text{Pl}\omega\text{NK}$, which additionally supports λ -abstractions, we instead encode the network topology (which features a lot of repetition) more concisely.⁴ First we identify several recurring patterns and give them appropriate names. These are shown in the top-right corner of Figure 1. The primitive patterns are sending and receiving, which are combined to create a link. Functions such as *send*, *recv* and *link* could be defined in a library or a language prelude, to be reused by other programs.

We can now re-write the topology as show on the right of Figure 1. The links from 1 to 3, 2 to 4 and 3 to 4 simply call the appropriate function. The link from 1 to 2 must directly rely on sending and receiving, but even here, we can see the benefits of the approach in reducing duplication.

The functional re-write of the routing program p features a more advanced use of functions (program p on the right-hand side of Figure 1). The function *forward* takes a source node (*src*) and a forwarding action *act* to perform. This action is a *thunk*, a suspended computation, which can be executed or *forced*, with the primitive *force*. The justification for these constructions is explained in the next section. Since thunks are essentially functions, *forward* is a higher-order function. A more extensive use of higher-order functions can be found in the Gossip Protocols example provided with the supplementary material.⁵

⁴The code presented here is untyped, for didactic purposes. From Section 3 onwards we will use typed $\text{Pl}\omega\text{NK}$, although type reconstruction for $\text{Pl}\omega\text{NK}$ is not difficult.

⁵<https://bitbucket.org/AlexanderV/probnetkat-lambda/src/1d12d/gossip-protocols.pnk>

To create the action *act* we use the function *to*. This function sets the packet's *pt* header to the given destination *dst*.

In *p*, the forwarding program for node 1 calls *to*, immediately suspends the call (using *thunk*), and passes it to *forward*. Forwarding for node 2 proceeds in a similar fashion.

The main expression is as before. Due to the additional structure, the readability of the program has improved considerably. Furthermore, code re-use has gone up, making the program easier to change. For instance, to model an additional link, from 1 to 4, we only need to add the call (*link* 1 4), instead of the more lengthy (*sw* = 1 ; *pt* = 4 ; *sw* $\leftarrow$ 4 ; *pt* $\leftarrow$ 1).

For another example, suppose we want to change the forwarding behaviour of our network, such that node 1 now chooses to forward to either node 2 or node 3 with equal probability, then we need only extend *p* slightly, obtaining *p'*. The necessary additions have been highlighted in Figure 1. In short, this section demonstrates the advantage for readability and maintainability that Pl ω NK provides.

2.4 Explicit Thunks and Sequencing

As we have shown in the previous section, it is highly desirable that functions are higher-order, in the sense that functions—and PNK expressions—can occur as arguments to other functions (e.g., *forward*). Then it seems reasonable to expect to be able to write the following:

$$sw \leftarrow 0; (\lambda x. sw = 1) (sw \leftarrow 1)$$

However, this presents an issue, since *sw* $\leftarrow$ 1 has a side-effect: it sets a header in the packet. The evaluation order is now important: if *sw* $\leftarrow$ 1 is evaluated before the application, this program accepts the packet, otherwise it drops the packet. The former corresponds to a Call-By-Value (CBV) order, the latter to a Call-By-Name (CBN) order. There is no clear reason to prefer one over the other, and both are useful in practice.

Indeed, we decide to not fix any particular evaluation order. Instead, we segregate terms into values and computations, inspired by the syntax of Call-By-Push-Value (CBPV) [Levy 2001]. Computations can be evaluated (possibly with side-effects), as opposed to values, which cannot be directly evaluated. Functions (classified as computations themselves) can only be applied to values. PNK terms are also computations, so the expression above is invalid in Pl ω NK syntax, since *sw* $\leftarrow$ 1 is not a value.

Instead, we obtain two possible variants, depending on whether CBV or CBN is intended (here the *to* is a primitive, *not* the function *to* defined earlier):

$$\text{CBV: } sw \leftarrow 0; (sw \leftarrow 1) \text{ to } y. ((\lambda x. sw = 1) y)$$

$$\text{CBN: } sw \leftarrow 0; (\lambda x. sw = 1) (\text{thunk } (sw \leftarrow 1))$$

In the first case, the primitive *to* (sequencing) evaluates *sw* $\leftarrow$ 1, including side effects, and binds the value that is produced to the variable *y*, followed by applying the function to *y*. In the second case, *sw* $\leftarrow$ 1 is thunked, and the function is applied to this thunk instead.

2.5 Semantics of Iteration

Foster et al. [2016] define the semantics of PNK in terms of measure theory. Semantically, a program denotes a function that maps sets of packet histories to a probability distribution over sets of packet histories. For example, the program *src* = 1, given input set *A*, returns a probability distribution that has probability 1 at the set $\{\pi :: h \in A \mid \pi.sw = 1\}$ (the notation $\pi.sw$ refers to the value of the *sw* header of π) and is zero everywhere else.

Iteration (*) is defined as an infinite stochastic process. The formalisation of this process is quite involved. Smolka et al. [2017b] give an equivalent, but much simpler definition, based on standard notions from domain theory.

In our work we retain this much simpler second definition, but extend it to support higher-order functions. However, since their domain is measure-theoretic, it does not support such functions. For this reason, we cannot use their domain directly. Instead, we rely on ω QBS, a domain developed by Vákár et al. [2019], which has domain-theoretic structure, supports measure-theory-like operations and admits higher-order functions.

2.6 Key Ideas

Denotational Semantics. We give a denotational semantics to our language within the ω QBS framework. In addition to PNK's semantics, ours also manages variable environments, and passes values (e.g. headers, constants, thunks) instead of just sets of packet histories.

The well-definedness of the semantics depends on two properties: First, to apply the domain-theoretic approach, we need to show that a specific notion of *continuity* holds. Second, we need a property that is similar to, or a consequence of, *strong normalisation*, but for denotational semantics. Informally, this property says that the denotation of a program produces only finitely many distinct values in a parallel fashion. The proofs for both properties have a similar structure: they proceed by induction on typing derivations of $\text{Pl}\omega\text{NK}$ programs and make use of logical relations. The definitions of these logical relations are interesting in their own right (see Sections 5.3 and 5.4).

Approximation and Decidable Equivalence. The streamlined semantics of Smolka et al. [2017b] formalises an iterative approximation procedure for PNK programs. The idea is to expand the iterations (*) up to n times, for some finite n . We define a procedure to compile a $\text{Pl}\omega\text{NK}$ program to PNK, while preserving the denotational semantics. Because our semantics is conservative with respect to the semantics of PNK, we can approximate the compiled program. Unfortunately, this compilation is only valid for a subclass of $\text{Pl}\omega\text{NK}$ programs. Essentially, the trick is to impose additional restrictions in the type system, such that the parallelism in $\text{Pl}\omega\text{NK}$ is limited to the parallelism that occurs in PNK.

Moreover, without the *dup* operation, PNK exhibits decidable program equivalence [Smolka et al. 2017a]. The *dup* operation duplicates the packet at the head of a packet history. Thus, removing this operation restrains all packet histories to the same length, making the state space of a program essentially discrete and finite. We conjecture that the same restriction also makes $\text{Pl}\omega\text{NK}$'s state space discrete. However, under this restriction, PNK produces only discrete distributions. Many useful properties are expressible in this sub-language [Smolka et al. 2019]. However, it cannot express some relevant properties, e.g. latency. Hence, this setting is less interesting than the full language. For this reason, we focus on full $\text{Pl}\omega\text{NK}$. We revisit these issues in Section 6.

3 SYNTAX AND TYPE SYSTEM

3.1 Syntax of Terms

The syntax of $\text{Pl}\omega\text{NK}$ (Figure 2) is a straightforward extension of the syntax of PNK with higher-order functions, thunks and sequencing. We segregate syntax terms into values and computations.

Values V are either variables x , *unit* values, header names h (from a finite set *Headers*, (e.g., *sw*, *pt*), literals n (natural numbers, the values that can be assigned to a header) or *thunks* (suspended computations). Values are never evaluated, but thunks can be forced.

Computations C , on the other hand, can be evaluated, but cannot occur as the argument to a function. Consequently, only values can appear on the right-hand side of an application.

Terms.

$x, y \in \text{Var}$	Variables
$h \in \text{Headers}$	Header names
$n \in \mathbb{N}$	Header values
$r \in \mathbb{R}$	Weights
$V = x \mid \text{unit} \mid h \mid n \mid \text{thunk } C$	Values
$P = \text{skip} \mid \text{drop} \mid V = V \mid \neg P \mid P \wedge P \mid P \vee P$	Predicates
$C = P \mid V \leftarrow V \mid \text{dup} \mid C; C \mid C \& C \mid C \oplus_r C \mid C^*$	PNK computations
$\mid \text{produce } V \mid \text{force } V \mid C \text{ to } x.C \mid \lambda x:S.C \mid C V$	New computations

Types and Contexts.

$S = \mathbf{1} \mid \mathcal{H} \mid \mathbb{N} \mid \mathcal{T} T$	Value types
$T = S \rightarrow T \mid \mathcal{P} S$	Computation types
$\Gamma = \emptyset \mid x:S, \Gamma$	Contexts

Fig. 2. Pl ω NK syntax.

All original PNK constructs are computations C . *Predicates* P are a subsort of those. *Atomic* predicates are *skip*, *drop* or tests $V = V$. *Composite* predicates are negation $\neg$, conjunction ($\wedge$) or disjunction ($\vee$).⁶ The remaining features inherited from PNK are (non-predicate) computations in Pl ω NK. They assign values to headers $V \leftarrow V$, duplicate packets with *dup*, compose sequentially ($;$) or in parallel ($\&$), make a probabilistic choice $\oplus_r$ (we sometimes elide the weight r when it is not relevant) or iterate $*$. We call the computations that Pl ω NK inherits from PNK the *probabilistic* computations.

Finally, Pl ω NK adds the following computations to PNK: producing a value V (*produce* V), forcing a thunk (*force* V), sequencing computations with C_1 *to* $x.C_2$, defining a function $\lambda x:S.C$, or applying functions C to values V with $(C V)$.

We also define *terminal* computations, i.e., computations that cannot be further evaluated:

$$R = P \mid V \leftarrow V \mid \text{dup} \mid R; R \mid R \& R \mid R \oplus_r R \mid R^* \\ \mid \text{produce } V \mid \lambda x:T.C$$

3.2 Types and Type System

Because PNK terms cannot “go wrong” or get stuck, the language did not come with a type system. This is no longer true for Pl ω NK, which introduces stuck terms with the lambda calculus. For this reason, we enrich Pl ω NK with a simple type system.

The chief reason for choosing *simple* types is strong-normalisation (see Theorem 6.8). In order to not compromise Pl ω NK’s suitability as a modelling language, certain properties, such as approximation and program equivalence must remain decidable, requiring terminating reduction for all well-typed programs.

For simplicity, we elide other less essential features, such as sum and product types, but such features could be added without too much trouble.

⁶Contrary to previous work [Foster et al. 2016], we do distinguish the syntax for disjunctive and conjunctive predicates from parallel and sequential composition of computations, for improved clarity. Earlier developments used the same operators for both.

$$\begin{array}{c}
\boxed{\Gamma \vdash_v V : S} \\
\\
\frac{x : S \in \Gamma}{\Gamma \vdash_v x : S} \quad \frac{}{\Gamma \vdash_v \text{unit} : \mathbf{1}} \quad \frac{}{\Gamma \vdash_v h : \mathcal{H}} \quad \frac{}{\Gamma \vdash_v n : \mathbb{N}} \quad \frac{\Gamma \vdash_c C : T}{\Gamma \vdash_v \text{thunk } C : \mathcal{T} T} \\
\boxed{\Gamma \vdash_c C : T} \\
\\
\frac{}{\Gamma \vdash_c \text{skip} : \mathcal{P} \mathbf{1}} \quad \frac{}{\Gamma \vdash_c \text{drop} : \mathcal{P} \mathbf{1}} \quad \frac{\Gamma \vdash_v V_1 : \mathcal{H} \quad \Gamma \vdash_v V_2 : \mathbb{N}}{\Gamma \vdash_c V_1 = V_2 : \mathcal{P} \mathbf{1}} \quad \frac{\Gamma \vdash_c P : \mathcal{P} \mathbf{1}}{\Gamma \vdash_c \neg P : \mathcal{P} \mathbf{1}} \\
\\
\frac{\Gamma \vdash_c P_1 : \mathcal{P} \mathbf{1} \quad \Gamma \vdash_c P_2 : \mathcal{P} \mathbf{1}}{\Gamma \vdash_c P_1 \wedge P_2 : \mathcal{P} \mathbf{1}} \quad \frac{\Gamma \vdash_c P_1 : \mathcal{P} \mathbf{1} \quad \Gamma \vdash_c P_2 : \mathcal{P} \mathbf{1}}{\Gamma \vdash_c P_1 \vee P_2 : \mathcal{P} \mathbf{1}} \quad \frac{\Gamma \vdash_v V_1 : \mathcal{H} \quad \Gamma \vdash_v V_2 : \mathbb{N}}{\Gamma \vdash_c V_1 \leftarrow V_2 : \mathcal{P} \mathbf{1}} \\
\\
\frac{}{\Gamma \vdash_c \text{dup} : \mathcal{P} \mathbf{1}} \quad \frac{\Gamma \vdash_c C_1 : T_1 \quad \Gamma \vdash_c C_2 : T_2}{\Gamma \vdash_c C_1 ; C_2 : T_2} \quad \frac{\Gamma \vdash_c C_1 : T \quad \Gamma \vdash_c C_2 : T}{\Gamma \vdash_c C_1 \& C_2 : T} \\
\\
\frac{\Gamma \vdash_c C_1 : T \quad \Gamma \vdash_c C_2 : T}{\Gamma \vdash_c C_1 \oplus C_2 : T} \quad \frac{\Gamma \vdash_c C : \mathcal{P} \mathbf{1}}{\Gamma \vdash_c C^* : \mathcal{P} \mathbf{1}} \quad \frac{\Gamma \vdash_v V : S}{\Gamma \vdash_c \text{produce } V : \mathcal{P} S} \quad \frac{\Gamma \vdash_v V : \mathcal{T} T}{\Gamma \vdash_c \text{force } V : T} \\
\\
\frac{\Gamma \vdash_c C_1 : \mathcal{P} S \quad x : S, \Gamma \vdash_c C_2 : T}{\Gamma \vdash_c C_1 \text{ to } x.C_2 : T} \quad \frac{x : S \Gamma \vdash_c C : T}{\Gamma \vdash_c \lambda x : S.C : S \rightarrow T} \quad \frac{\Gamma \vdash_c C : S \rightarrow T \quad \Gamma \vdash_v V : S}{\Gamma \vdash_c C V : T}
\end{array}$$

Fig. 3. $\text{Pl}\omega\text{NK}$'s typing rules.

The bottom part of Figure 2 shows the syntax of types. Types are divided into two kinds: value types and computation types. There are 4 forms of value types: unit types $\mathbf{1}$, header labels $\mathcal{H}$, header literals $\mathbb{N}$ and thanks $\mathcal{T} T$. Furthermore, there are 2 forms of computation types: function types $S \rightarrow T$ and producer types $\mathcal{P} S$. By construction, the argument position of a function type can only be a value type. Likewise, only value types can appear inside producer ($\mathcal{P}$) types. Since only these constructions bind variables, only value types appear in contexts.

Figure 3 shows the typing rules. We define two mutually recursive typing judgements: $\Gamma \vdash_v V : S$ for typing value terms, and $\Gamma \vdash_c C : T$ for typing computation terms.

The rules for values are straightforward. Predicates are typed as computations, and always have type $\mathcal{P} \mathbf{1}$ because they do not produce useful results, only useful side-effects. This also applies to the rules for modification and duplication (see below).

When sequentially composing computations $C_1 ; C_2$, only C_2 determines the type of the whole computation. For parallel composition and choice, the types of both computations must be *the same*. Intuitively, the former discards the value produced by the first computation, while the latter two must somehow combine the two produced values, requiring them to have the same type.

Iteration C^* has the same type as C , which is only allowed to be $\mathcal{P} \mathbf{1}$, for the following reason. Consider that the iteration could be zero or more times, and thus in the zero case would require inventing a value of an arbitrary type, which we cannot do unless we restrict iteration to a fixed type with a known value—the unit type. The remaining rules in Figure 3 are derived from CBPV [Levy 2001]. However, unlike CBPV, $\text{Pl}\omega\text{NK}$ does *not* exhibit certain type isomorphisms, for instance:

$$S \rightarrow S' \rightarrow T \not\cong S' \rightarrow S \rightarrow T$$

This is because in our language computations of function type can have side-effects without being applied, whereas in CBPV computation of function type are only evaluated upon application. In particular, applying a computation of type $S \rightarrow S' \rightarrow T$ to a value x produces a computation of type $S' \rightarrow T$ which may have side-effects that depend on x .

In the previous sections, we have also used “top-level” definitions of the form $f = \dots$. These are to be understood as syntactic sugar. They desugar as follows:

$$\boxed{\begin{array}{l} f_1 = C_1 \\ C_2 \end{array}} \rightsquigarrow (\lambda y. \mathcal{T} T. [f_1 \mapsto \text{force } y] C_2) \text{ (thunk } C_1 \text{)}$$

where y is fresh, and $\vdash_c C_1 : T$. Recall that only values may be bound to variables. Since the C_1 on the right-hand side of the equality is a computation, we first convert this expressions to a value by thunking them, and then forcing them where they occur in C_2 .

4 A CONVENIENT CATEGORY FOR $\text{Pl}\omega\text{NK}$

Measure theory is the usual model for continuous⁷ distributions. However, for our case classical measure theory has a critical shortcoming: function spaces of measure spaces are not necessarily measurable themselves, making measure theory unsuitable as the model of a programming language with λ -abstraction [Aumann et al. 1961]. Put in a different way, the category of measurable spaces is not Cartesian closed.

Quasi-Borel Spaces [Heunen et al. 2017] are a recent advancement in the state-of-the-art of the semantics of probabilistic programs, which *are* Cartesian closed and provide an alternative formalisation of probabilistic structures. An even more recent development are the ω -Quasi-Borel Spaces [Vákár et al. 2019], which additionally provide ωCPO structure. We require this structure to model iterations.

The remainder of this section provides a brief overview of ω -Quasi-Borel Spaces, which we use as the semantic domain for the denotational semantics of $\text{Pl}\omega\text{NK}$ that is presented in the next section. Eager readers may skip ahead to Section 5 on a first reading and come back when they want more detail.

4.1 ω -Complete Partial Orders

Semantically modelling the behaviour of $\text{Pl}\omega\text{NK}$ or indeed plain PNK requires a semantic domain that captures the recursive nature of iterations C^* . For this purpose we use a partially ordered domain and make sure that the denotation increases with each iteration step, i.e., it ascends. The result of the entire iteration is then given by the least upper bound of the denotations of the iteration steps, and we must choose the ordering in such a way that these least upper bounds always exist. The ω -Complete Partial Orders (ωCPOs) are the orders with this property: in an ωCPO the least upper bounds of ascending chains always exist. Let us now define these concepts more precisely.

Definition 4.1. Let $(P, \sqsubseteq)$ be some poset, an ω -chain is a sequence $(x_n)_{n \in \mathbb{N}}$ for $x_n \in P$, such that $\forall i, j \in \mathbb{N} : i \leq j$ implies $x_i \sqsubseteq x_j$. We will sometimes write such a chain as $x_0 \sqsubseteq x_1 \sqsubseteq \dots$.

Definition 4.2. The poset P is an ω -Complete Partial Order (ωCPO) when every ω -chain has a least upper bound (lub) $\bigsqcup_{n \geq 0} x_n \in P$ (sometimes also denoted $\bigvee_{n \geq 0} x_n$). The *least upper bound* is

⁷At first glance, it might seem counter-intuitive that $\text{Pl}\omega\text{NK}$ could admit continuous probability distributions. However, Foster et al. [Foster et al. 2016] show an example of a PNK program where this is the case: let p be a program that outputs two distinct packets with equal probability, then $p ; (\text{dup} ; p)^*$ denotes a continuous distribution.

the smallest element of P that is larger than every x_n . More formally,

$$\forall n \in \mathbb{N} : x_n \sqsubseteq \bigsqcup_{n \geq 0} x_n, \text{ and } \forall z \in P : (\forall n \in \mathbb{N} : x_n \sqsubseteq z) \Rightarrow \bigsqcup_{n \geq 0} x_n \sqsubseteq z.$$

Example 4.3. The powerset 2^X of any set X is an ω CPO when ordered by subset inclusion ($\subseteq$). The lub of any ω -chain $X_0 \subseteq X_1 \subseteq \dots$ (with $X_i \subseteq X$) is precisely the union $\bigcup_{i \geq 0} X_i$. In fact, $\langle 2^X, \subseteq \rangle$ is a complete lattice, meaning that any subset of 2^X has a lub.

Example 4.4. The functions $f : X \rightarrow P$ into an ω CPO $\langle P, \sqsubseteq_P \rangle$ also form an ω CPO, under the pointwise order $\leq$, defined as $f \leq g \iff \forall x \in X : f(x) \sqsubseteq_P g(x)$. In this instance, we usually denote the lub by $\bigvee$.

Continuous functions between two ω CPOs $\langle P, \sqsubseteq_P \rangle$ and $\langle Q, \sqsubseteq_Q \rangle$ are monotone functions $f : P \rightarrow Q$ such that f preserves least upper bounds, i.e. for all ω -chains $(a_n)_{n \in \mathbb{N}}$ in P ,

$$f\left(\bigsqcup_{n \geq 0} a_n\right) = \bigsqcup_{n \geq 0} f(a_n).$$

Note that the first lub is in P , the second in Q . The monotonicity requirement ensures that this second lub actually exists, by ensuring that $(f(a_n))_{n \in \mathbb{N}}$ is an ω -chain.

The ω CPOs and continuous functions between them form a Cartesian closed category, i.e.:

- the composition of two continuous functions is continuous;
- ω CPOs are closed under Cartesian products, i.e., given two ω CPOs P and Q , $P \times Q$ is also an ω CPO; and
- continuous functions between two ω CPOs also form an ω CPO. The order is the pointwise order.

4.2 Quasi-Borel Spaces

In $\text{Pl}\omega\text{NK}$ we give semantics to higher-order functions and probabilistic choice. For this reason we need a semantic domain which is both Cartesian closed and admits a probabilistic power domain. As we mentioned previously, we cannot rely on measure theory, the usual choice for probabilistic domains, since it is not Cartesian closed. Instead, we use Quasi-Borel Spaces (QBS) [Heunen et al. 2017], which are Cartesian closed.

However, in order to define QBSes, we must revisit some basic definitions from measure theory first. We limit the treatment of measure theory to the essentials needed to understand Quasi-Borel Spaces. For instance, we do not discuss general measure spaces, but constrain ourselves to a particular set of measurable sets on the reals, the Borel sets:

Definition 4.5. The Borel sets $\mathfrak{B}$ are the least collection of subsets of $\mathbb{R}$, such that:

- the intervals $[a, b]$ are Borel sets for $a, b \in \mathbb{R}$,
- the complement of a Borel set is a Borel set, and
- countable unions of Borel sets are Borel sets.

A probability measure is a function $\mu : \mathfrak{B} \rightarrow [0, 1]$ satisfying $\mu(\mathbb{R}) = 1$ and $\mu(\bigcup S_n) = \sum \mu(U_n)$, for any countable sequence of disjoint Borel sets $(S_n)_{n \in \mathbb{N}}$. A function $f : \mathbb{R} \rightarrow \mathbb{R}$ is called *measurable* if its inverse image maps Borel sets to Borel sets. Symbolically, for all $B \in \mathfrak{B}$:

$$f^{-1}(B) = \{x \in \mathbb{R} \mid f(x) \in B\} \in \mathfrak{B}.$$

Such functions can be integrated with respect to a measure. For a *non-negative* real-valued measurable function $f : \mathbb{R} \rightarrow \mathbb{R}$, the integral of f with respect to a probability measure μ is defined

as:

$$\int f \, d\mu = \sup_{(U_n)} \sum_n \left(\mu(U_n) \inf_{x \in U_n} f(x) \right)$$

where (U_n) ranges over finite partitionings of $\mathbb{R}$ into Borel sets. If f is allowed to be negative, its integral is:

$$\int f \, d\mu = \int f^+ \, d\mu - \int f^- \, d\mu$$

where $f^+ = \max(f, 0)$, $f^- = \max(-f, 0)$.

Measure theory generalises Borel sets to the measurable subsets of a measurable space X , and generalises measurable functions on $\mathbb{R}$ to measurable functions between measurable spaces, whose inverse image always maps measurable sets to measurable sets.

Quasi-Borel Spaces [Heunen et al. 2017] are an alternative to measure theory, starting from the from the notion of a random element $\mathbb{R} \rightarrow X$ instead of measurable sets.

Definition 4.6. A Quasi-Borel Space (QBS) $\langle X, M_X \rangle$ is a set X together with a set of functions M_X such that the following conditions are met:

- If $\alpha : \mathbb{R} \rightarrow X$ is constant, then $\alpha \in M_X$;
- If $\alpha \in M_X$, and $f : \mathbb{R} \rightarrow \mathbb{R}$ is measurable, then $\alpha \circ f \in M_X$;
- let $\mathbb{R} = \bigcup_{n \in \mathbb{N}} U_n$ where the U_n are pairwise disjoint Borel sets, if $\alpha_n \in M_X$ for all $n \in \mathbb{N}$, then $\beta \in M_X$ where $\beta(r) = \alpha_n(r)$ if $r \in U_n$.

Essentially, M_X must contain all constant functions, and must be closed under pre-composition with a measurable function or countable case-splitting. A function $f : X \rightarrow Y$ is a morphism from $\langle X, M_X \rangle$ to $\langle Y, M_Y \rangle$ if for all $\alpha \in M_X$, $f \circ \alpha \in M_Y$.

There are two canonical ways to turn a set X into a QBS. One option is to simply include *all* functions $\mathbb{R} \rightarrow X$ in M_X . Another option is to take as random elements all measurably piece-wise constant functions, i.e., those functions from $\mathbb{R}$ to X that are piece-wise constant on Borel sets.

The quasi-Borel spaces together with their morphisms form a category. Moreover, this category admits products, co-products and function spaces. That is, unlike the category formed by measurable spaces and measurable functions between them, this category is Cartesian closed [Heunen et al. 2017, Proposition 18].

4.3 ω -Quasi-Borel Spaces

Definition 4.7. An ω -Quasi-Borel Space is a triple $\langle X, M_X, \sqsubseteq_X \rangle$ such that:

- $\langle X, M_X \rangle$ is a quasi-Borel Space,
- $\langle X, \sqsubseteq_X \rangle$ is an ω CPO, and
- $(\bigvee_{n \geq 0} \alpha_n) \in M_X$ for all ω -chains $\alpha_0 \leq \alpha_1 \leq \dots$ (with $\alpha_n \in M_X$), where $\leq$ is the point-wise order on M_X .

Informally, an ω QBS is a QBS that is also an ω CPO, and whose random elements are closed under pointwise lubs of ω -chains.

The morphisms between ω QBSes are those morphisms between the underlying ω QBSes, which are also Scott-continuous between the underlying ω CPOs.

We reiterate two examples from Vákár et al. [2019]:

Example 4.8. Real values have the ω QBS, $\langle \mathbb{R}, M_{\mathbb{R}}, = \rangle$, with $M_{\mathbb{R}}$ the set of measurable functions from $\mathbb{R}$ to $\mathbb{R}$. Alternatively, consider $\mathbb{W} = \langle [0, \infty], M_{\mathbb{W}}, < \rangle$ the space of weights, where $M_{\mathbb{W}}$ is the set of measurable functions from $\mathbb{R}$ to $[0, \infty]$.

$$\begin{aligned}
\llbracket \mathbf{1} \rrbracket &= \langle \{()\}, M_{\{()\}}, (=) \rangle \\
\llbracket \mathcal{H} \rrbracket &= \langle \text{Headers}, M_{\text{Headers}}, (=) \rangle \\
\llbracket \mathbb{N} \rrbracket &= \langle \mathbb{N}, M_{\mathbb{N}}, (=) \rangle \\
\llbracket \mathcal{T} T \rrbracket &= \llbracket T \rrbracket \\
\llbracket S \rightarrow T \rrbracket &= 2^{PH} \rightarrow D(\llbracket T \rrbracket^{[S]} \rightarrow 2^{PH}) \\
\llbracket \mathcal{P} S \rrbracket &= 2^{PH} \rightarrow D(\llbracket S \rrbracket \rightarrow 2^{PH}) \\
A \multimap B &= A \rightarrow B_{\perp}
\end{aligned}$$

Fig. 4. Denotations of types.

Just like QBS, ω QBS forms a category that is closed under products, co-products and exponentials (functions). This category is Cartesian closed. With ω QBSes we finally have a mathematical concept which unifies iteration, probabilistic choice and higher-order functions.

4.4 A Commutative Probabilistic Powerdomain

The denotational semantics we give to $\text{Pl}\omega\text{NK}$ is a monadic semantics: it allows the structuring of the semantics in a compositional fashion [Moggi 1991]. This section explains how to define a monad suitable for expressing probabilistic computations as an ω QBS.

The idea, as explained by Vákár et al. [2019], is to treat distributions as expectation or integration operators. The distribution monad $D(X)$ is (a submonad⁸ of) the continuation monad $C(X) = (X \rightarrow \mathbb{W}) \rightarrow \mathbb{W}$, where the arrows ($\rightarrow$) denote ω QBS morphisms. The idea is that for a $\mu \in C(X)$ and $f : X \rightarrow \mathbb{W}$, $\mu(f)$ computes the integral of f with respect to μ , in more traditional notation, $\mu(f) = \int f \, d\mu$.

The unit (*return*) and composition ($\gg$) of this monad are given by:

$$\text{return } x = \lambda k \rightarrow k \, x \qquad m \gg f = \lambda k \rightarrow m \, (\lambda x \rightarrow f \, x \, k)$$

In more traditional language these represent integrating with δ_x (Dirac-delta of x) and the integral $\int_x \int k \, d(f \, x) \, dm$, respectively. The details of this construction are beyond the scope of this article. The most important results are:

- This construction satisfies the requirements of synthetic measure theory, implying that the results of measure theory continue to hold in this new setting. In particular, we can do addition and scalar multiplication of distributions.
- The monad consists of the *s-finite measures and kernels*, meaning that the monad is commutative: operations can be re-ordered [Staton 2017].

5 DENOTATIONAL SEMANTICS

A denotational semantics maps terms (values and computations) into a semantic domain. A term's domain is determined by its type, therefore, we must foremost discuss the semantics of types.

5.1 Semantics of Types

Figure 4 shows the interpretations of the different types. Each type denotes an ω QBS, which we denote by just the underlying set X instead of the full triple $\langle X, M_X, \sqsubseteq_X \rangle$.

⁸Specifically, it is the smallest ω QBS that is a full sub- ω CPO of $C(X)$ and contains the randomisable random elements.

Unit types **1** are denoted by nullary products, whose only inhabitant is written as $()$. Header types $\mathcal{H}$ are denoted by the finite set of header labels and header literals $\mathbb{N}$ are denoted by the natural numbers. In these three cases, the underlying ω CPO is given by the discrete order $(=)$ ⁹ and the underlying ω QBS can be created using either of the canonical methods described in Section 4.2. Thunk types simply denote the denotation of the thunked computation type.

Computation types need to capture the side effects: state, parallelism, and probabilistic choice. The denotations of both sorts of computation types are of a similar form:

$$\begin{array}{ccc} \text{input state} & & \text{output state} \\ \underbrace{2^{PH}} & \rightarrow D(\underbrace{X \multimap 2^{PH}}_{\text{parallel value and state}}) \end{array}$$

where D is the distribution monad defined in Section 4.4, the arrow $(\rightarrow)$ denotes ω QBS-morphism and the harpoon $(\multimap)$ denotes partial maps. Partial maps $f : A \multimap B$ are equivalent to total functions $\widehat{f} : A \rightarrow B_\perp$, where $B_\perp$ is B extended with a distinguished element, $\perp_B$, preceding all other elements. For any $a \in A$, $\widehat{f}(a) = \perp_B$ then means that f is undefined on a . Hence, the domain $\text{dom}(f)$ of a partial map f is defined as $\text{dom}(f) = \{x \mid \widehat{f}(x) \neq \perp\}$. The partial maps are not required to be continuous (but they do form an ω CPO and an ω QBS).

Note that we use the power set of packet histories to represent the state. The ω QBS of this powerset is $\langle 2^{PH}, M_{2^{PH}}, \subseteq \rangle$ where $M_{2^{PH}}$ are the random elements of the exponential 2^{PH} of $\langle 2, M_2, = \rangle$ and $\langle PH, M_{PH}, = \rangle$.

We use partial maps $X \multimap 2^{PH}$ to model parallelism in both the value (X) and the state (2^{PH}). The idea is that a particular value comes with a particular state. Since the map is partial, not all values are necessarily present. Note that we cannot use total maps and the empty set to indicate the absence of a value, since the empty set is a *valid* state.

For producer types $\mathcal{P} S$, $X = \llbracket S \rrbracket$. For function types $S \rightarrow T$, $X = \llbracket T \rrbracket^{[S]}$, i.e., the morphisms (functions) from $\llbracket S \rrbracket$ to $\llbracket T \rrbracket$.

Finally, the semantics $\llbracket S \rrbracket$ extends naturally to an interpretation on contexts (recall that contexts only contain value types):

$$\llbracket x_1 : S_1, \dots, x_k : S_k \rrbracket = \llbracket S_1 \rrbracket \times \dots \times \llbracket S_k \rrbracket$$

That is, contexts are denoted by the product of the denotations of the types within the context. For convenience, however, for an environment $\rho \in \llbracket \Gamma \rrbracket$, we will write $\rho(x)$ to look up a variable x , and $[x \mapsto v]\rho$ to update the value of x in ρ .

5.2 Semantics of Terms

Next, we define denotational semantics for terms (Figure 5). The semantics are divided into three (mutually recursive) categories: semantics for values, for predicates and for other computations. All semantic functions take an environment $\rho \in \llbracket \Gamma \rrbracket$ as their first argument. Predicates and computations also take a set of packet histories $A \in 2^{PH}$.

5.2.1 Values. For values, the semantics is relatively straightforward: variables can be looked up in the environment ρ . The *units*, headers and literals map to their corresponding constants. A thunk *thunk* C partially applies the denotation of C to the current environment ρ . Recall that $\llbracket C \rrbracket \rho$ is a function that expects a set of packet histories. When the thunk is forced this function is applied to the current state. To see why this makes sense, consider that it is the location where the thunk

⁹In this order, everything is incomparable to everything, unless they are the same element.

Values

$$\llbracket V \rrbracket : \llbracket \Gamma \rrbracket \rightarrow \llbracket S \rrbracket \text{ for } \Gamma \vdash_v V : S$$

$$\llbracket x \rrbracket \rho = \rho(x)$$

$$\llbracket \text{unit} \rrbracket \rho = ()$$

$$\llbracket h_i \rrbracket \rho = h_i$$

$$\llbracket i \rrbracket \rho = i$$

$$\llbracket \text{thunk } C \rrbracket \rho = \llbracket C \rrbracket \rho$$
Predicates $\llbracket P \rrbracket^P : \llbracket \Gamma \rrbracket \rightarrow 2^{PH} \rightarrow 2^{PH}$ for $\Gamma \vdash_c P : \mathcal{P} \mathbf{1}$

$$\llbracket \text{skip} \rrbracket^P \rho A = A$$

$$\llbracket \text{drop} \rrbracket^P \rho A = \emptyset$$

$$\llbracket V_1 = V_2 \rrbracket^P \rho A = \{\pi :: h \in A \mid \pi.(\llbracket V_1 \rrbracket \rho) = \llbracket V_2 \rrbracket \rho\}$$

$$\llbracket \neg P \rrbracket^P \rho A = A - \llbracket P \rrbracket^P \rho A$$

$$\llbracket P_1 \wedge P_2 \rrbracket^P \rho A = \llbracket P_1 \rrbracket^P \rho A \cap \llbracket P_2 \rrbracket^P \rho A$$

$$\llbracket P_1 \vee P_2 \rrbracket^P \rho A = \llbracket P_1 \rrbracket^P \rho A \cup \llbracket P_2 \rrbracket^P \rho A$$
Computations $\llbracket C \rrbracket : \llbracket \Gamma \rrbracket \rightarrow \llbracket T \rrbracket$ for $\Gamma \vdash_c C : T$

$$\llbracket P \rrbracket \rho A = \text{return } (\lambda() \rightarrow \llbracket P \rrbracket^P \rho A)$$

$$\llbracket V_1 \leftarrow V_2 \rrbracket \rho A = \text{let } f = \llbracket V_1 \rrbracket \rho, j = \llbracket V_2 \rrbracket \rho \\ \text{in return } (\lambda() \rightarrow \{\pi[f \mapsto j] :: h \mid \pi :: h \in A\})$$

$$\llbracket \text{dup} \rrbracket \rho A = \text{return } (\lambda() \rightarrow \{\pi :: \pi :: h \mid \pi :: h \in A\})$$

$$\llbracket C_1 ; C_2 \rrbracket \rho A = \text{do } \mu \leftarrow \llbracket C_1 \rrbracket \rho A \\ \llbracket C_2 \rrbracket \rho \left(\bigcup_{x \in \text{dom}(\mu)} \mu(x) \right)$$

$$\llbracket C_1 \& C_2 \rrbracket \rho A = \text{do } \mu_1 \leftarrow \llbracket C_1 \rrbracket \rho A \\ \mu_2 \leftarrow \llbracket C_2 \rrbracket \rho A \\ \text{return } (\mu_1 \vee \mu_2)$$

$$\llbracket C_1 \oplus_r C_2 \rrbracket \rho A = r(\llbracket C_1 \rrbracket \rho A) + (1-r)(\llbracket C_2 \rrbracket \rho A)$$

$$\llbracket C^* \rrbracket \rho A = \bigsqcup_{n \geq 0} (\llbracket C^n \rrbracket \rho A) \text{ where } \\ C^0 = \text{skip}, \\ C^{n+1} = \text{skip} \& (C ; C^n), n \geq 0$$

$$\llbracket \text{produce } V \rrbracket \rho A = \text{return } (\lambda(\llbracket V \rrbracket \rho) \rightarrow A)$$

$$\llbracket \text{force } V \rrbracket \rho A = \llbracket V \rrbracket \rho A$$

$$\llbracket \lambda x : S. C \rrbracket \rho A = \text{return } (\lambda(\lambda v \rightarrow \llbracket C \rrbracket ([x \mapsto v] \rho)) \rightarrow A)$$

$$\llbracket C V \rrbracket \rho A = \text{do } \mu \leftarrow \llbracket C \rrbracket \rho A \\ \exists \{f(\llbracket V \rrbracket \rho) \mu(f) \mid f \in \text{dom}(\mu)\}$$

$$\llbracket C_1 \text{ to } x. C_2 \rrbracket \rho A = \text{do } \mu \leftarrow \llbracket C_1 \rrbracket \rho A \\ \exists \{\llbracket C_2 \rrbracket [x \mapsto v] \rho \mu(v) \mid v \in \text{dom}(\mu)\}$$

$$\text{where } \exists \{m_1, \dots, m_k\} = \text{do } \mu_1 \leftarrow m_1 \\ \vdots \\ \mu_k \leftarrow m_k \\ \text{return } (\mu_1 \vee \dots \vee \mu_k)$$

Note: $X \rightarrow Y$ means ω QBS-morphisms from X to Y .

Fig. 5. Denotations of terms.

is created that determines the scope of the *variables* in C , but it is the location where it is forced which determines its *state*.

5.2.2 Predicates. Predicates filter sets of packet histories, that is, they allow or reject particular packet histories.

Predicates *skip* and *drop* allow, respectively disallow, all packet histories. A guard $V_1 = V_2$ only allows packet histories where the first packet's header $\llbracket V_1 \rrbracket \rho$ has the specified value $\llbracket V_2 \rrbracket \rho$ (accessing header f of a packet π is written as $\pi.f$). Negation ($\neg$) only retains packets that are dropped by its argument. Finally, conjunction ($\wedge$) and disjunction ($\vee$) take the intersection and union respectively.

5.2.3 Computations. The semantics for computations returns values in the monad D . Since D is a monad, we use the well-understood *do-notation* to construct monadic expressions, de-sugaring straightforwardly to monadic bind ($\gg$).

The first three cases are non-probabilistic. Moreover, they all produce the unit value $()$, they return a map from $()$ to a modified set of histories.¹⁰ Predicates P filter packet histories according to the predicate semantics $\llbracket P \rrbracket^P \rho A$. Modifications change the first packet's header for every packet history in the input set (setting a packet π 's header f to x is written $\pi[f \mapsto x]$). Similarly, *dup* duplicates the first packet.

Sequential composition $(C_1 ; C_2)$ first evaluates C_1 , and then C_2 (returning the value of C_2). Since the evaluation of C_2 does not depend on the value of C_1 , all packet histories of C_1 are simply aggregated and passed to C_2 .

Parallel composition $C_1 \& C_2$ combines the results (values and states) of both computations. Since the results of C_1 and C_2 are captured by the partial maps μ_1 and μ_2 , it must combine these maps. The pointwise lub $\mu_1 \vee \mu_2$ is exactly what is needed: the domain of the lub is the union of the domains of μ_1 and μ_2 , so it has all the values of both computations, and the sets of packet histories it maps a value x onto is the union of $\mu_1(x)$ and $\mu_2(x)$.

Choice $C_1 \oplus_p C_2$, reweighs both C_1 and C_2 , by p and $1 - p$ respectively, and adds the resulting distributions. Note that the sum is a probability distribution, i.e. all weights sum to one. Since D satisfies the requirements of synthetic measure theory, scalar multiplication and addition behave as one would intuitively expect, multiplying and adding probabilities.

Iteration (Kleene-star C^*) is defined as it is for PNK, using least fixed-point semantics. It can be understood as the least fixed-point of the function $(\gg \lambda \mu \rightarrow \llbracket \text{skip} \& C \rrbracket \rho \mu())$. Or, operationally, we simply take the least upper bound of iterating C for $0, 1, \dots$ times. This definition is identical to Smolka et al.'s [2017b], although on a different ω CPO. Treating probabilistic loops as fixed-points goes at least as far back as Kozen's [1981] early work on probabilistic program semantics.

To produce (*produce* V) a value V , we return the map from $\llbracket V \rrbracket \rho$ to the current set of packet histories A . Here we use the notation $\lambda(\llbracket V \rrbracket \rho) \rightarrow A$ to define the map that is A on $\llbracket V \rrbracket \rho$ and $\perp$ everywhere else.

Thunks are forced by evaluating their semantics, and then applying the result of that semantics to the input set of packet histories. Recall that the semantics of a thunk corresponds to a partially applied semantics of the thunked computation.

For an abstraction $\lambda x : S.C$ we construct a map with a domain containing only one particular anonymous function in the meta-language. This function takes a $v \in \llbracket S \rrbracket$, extends the environment ρ with v , and runs $\llbracket C \rrbracket$ in this new environment.

¹⁰By convention, λ with an arrow ($\rightarrow$) is abstraction in the meta-language (morphisms in ω QBS) and λ with dot $\cdot$ refers to abstraction in the object language.

For application, we must first sample from the computation C . This produces a map from $\llbracket S \rrbracket \rightarrow \llbracket T \rrbracket$ to sets of packet histories. Each *unique* function f in the map μ is then applied to the argument $\llbracket V \rrbracket \rho$ and the corresponding set of packet histories $\mu(f)$. Every unique function f corresponds to one or more parallel branches that produce this value, $\mu(f)$ aggregates the states of each of those branches. This produces a set of distributions, which we can collapse using the Ξ operator (also defined in Figure 5). This operator samples from each distribution, and takes the lub of the resulting maps.

Sequencing (C_1 to $x.C_2$) is similar to application. We sample a map μ from C_1 . As with application, there is only a finite number of distinct values v in $\text{dom}(\mu)$. For each unique v , we extend ρ , and evaluate $\llbracket C_2 \rrbracket [x \mapsto v] \rho \mu(v)$, producing a finite set of distributions, which we flatten with Ξ . The use of Ξ is well-defined by the following lemma:

LEMMA 5.1 (LUBS OF FINITE SETS). *Let $\{m_1, \dots, m_n\} \subseteq \llbracket T \rrbracket$ for some type T , then $\Xi\{m_1, \dots, m_n\}$ is the least upper bound of $\{m_1, \dots, m_n\}$.*

The proof is a straightforward induction on the size of the set. Because Ξ computes the least upper bound, it is independent of the order of $\{m_1, \dots, m_n\}$. However, it assumes that the input set is finite, i.e., that there are only finitely many unique functions f . This assumption is discharged by Theorem 5.2 (see Section 5.3).

A further theorem (Theorem 5.5) makes the semantics well-defined. In particular, it ensures that the least upper bounds in Figure 5 iteration exists and that the types ascribed to the denotations are valid. The next sections discuss each theorem in turn.

5.3 Finite Maps

The function Ξ is only defined for *finite* sets. Therefore, we must verify that this function is only ever applied to a finite set. Informally, this is the case if we ensure that in Figure 5, μ is only bound to maps that have a finite domain. More formally, we desire the following:

THEOREM 5.2 (FINITE MAPS). *For all computations C , such that $\Gamma \vdash_c C : T$, $\rho \in \llbracket \Gamma \rrbracket$ and $A \in 2^{PH}$, we have that for any ω QBS morphism $f : X \rightarrow \mathbb{R}$, where $\llbracket T \rrbracket = 2^{PH} \rightarrow D(X)$:*

$$\int_X f d(\llbracket C \rrbracket \rho A) = \int_X \chi_F f d(\llbracket C \rrbracket \rho A)$$

where $F = \{g \in X \mid \text{dom}(g) \text{ is finite}\}$, and $\chi_F : X \rightarrow \{0, 1\}$ is its characteristic function.

To prove this theorem we rely on logical relations [Tait 1967] $\mathcal{F}_S$ to define a stronger theorem (Theorem 5.4). The actual logical relations are $\mathcal{F}_S$ and $\mathcal{F}_T$, defined as follows:

Definition 5.3. The predicates $\mathcal{F}_S$ and $\mathcal{F}_T$ where S and T are value, respectively computations types, are defined inductively as:

$$\begin{aligned} \mathcal{F}_S &= \begin{cases} \llbracket S \rrbracket & \text{if } S \neq \mathcal{T} T \\ \mathcal{F}_T & \text{if } S = \mathcal{T} T \end{cases} \\ \mathcal{F}_T &= \left\{ f \in \llbracket T \rrbracket \mid \forall g \in G(T), \forall A \in 2^{PH} : \int g d(f(A)) = \int \chi_{F(T)} g d(f(A)) \right\} \end{aligned}$$

where

$$\begin{aligned} F(\mathcal{P} S) &= \{\mu : \mathcal{F}_S \rightarrow 2^{PH} \mid \text{dom}(\mu) \text{ is finite}\} \\ F(S \rightarrow T) &= \{\mu : (\mathcal{F}_S \rightarrow \mathcal{F}_T) \rightarrow 2^{PH} \mid \text{dom}(\mu) \text{ is finite}\} \\ G(\mathcal{P} S) &= (\llbracket S \rrbracket \rightarrow 2^{PH}) \rightarrow \mathbb{R} \\ G(S \rightarrow T) &= ((\llbracket S \rrbracket \rightarrow \llbracket T \rrbracket) \rightarrow 2^{PH}) \rightarrow \mathbb{R} \end{aligned}$$

The idea of these logical relations is to restrict the denotations to those semantic objects where all partial maps have a finite domain. For instance, the denotation of non-thunk values never contains a map, hence $\mathcal{F}_S$ is simply $\llbracket S \rrbracket$ in this case. Thunked computations can contain computations, thus we restrict them to $\mathcal{F}_T$.

In a somewhat roundabout fashion, $\mathcal{F}_T$ says that the denotations of a computation of type T must only contain finite maps. In particular, it says that integrating any g with respect to $f(A)$ should be the same as integrating $\chi_{F(T)} g$ with respect to $f(A)$, where $F(T)$ only contains the finite maps of the appropriate type.¹¹ In other words, integrating while filtering out the infinite maps should not make a difference at all, meaning that the only maps that are present are finite. The following theorem, of which Theorem 5.2 is a corollary, states that every well-typed term's denotation is a member of the appropriate relation:

THEOREM 5.4. *For all values V and computations C ,*

$$\text{if } \Gamma \vdash_v V : S \text{ and } \Gamma \vdash_c C : T, \text{ then } \llbracket V \rrbracket \rho \in \mathcal{F}_S \text{ and } \llbracket C \rrbracket \rho \in \mathcal{F}_T$$

where $\rho \in \llbracket \Gamma \rrbracket$ such that if $x : S_x \in \Gamma$, then $\rho(x) \in \mathcal{F}_{S_x}$.

PROOF. (sketch) By induction on the structure of the typing derivation, performing case analysis on the final rule application. The proof uses an equational reasoning style. $\square$

5.4 Continuity

Continuity is a property that intuitively means that a function preserves the least upper bounds of the domain it operates on. In other words, it preserves the ω CPO structure. This is a technical requirement to ensure that the least-upper bound of the semantics of iteration exists. More formally, we desire the following property:

THEOREM 5.5 (CONTINUITY). *For all computations C , such that $\Gamma \vdash_c C : T$, for all $\rho \in \llbracket \Gamma \rrbracket$, $\llbracket C \rrbracket \rho A$ is continuous in $A \in 2^{PH}$, i.e. $\llbracket C \rrbracket \rho A$ is monotone, and for all $A_1 \subseteq A_2 \subseteq \dots \subseteq PH$:*

$$\bigsqcup_{i \geq 1} (\llbracket C \rrbracket \rho A_i) = \llbracket C \rrbracket \rho \left(\bigcup_{i \geq 1} A_i \right)$$

To prove this theorem, we define the following logical relation:

Definition 5.6. The predicates C_S and C_T where S and T are value, respectively computations types, are defined inductively as:

$$\begin{aligned} C_S &= \begin{cases} \llbracket S \rrbracket & \text{if } S \neq \mathcal{T} T \\ C_T & \text{if } S = \mathcal{T} T \end{cases} \\ C_T &= \left\{ f \in \llbracket T \rrbracket \mid \begin{array}{l} f \text{ is continuous, and} \\ \forall g \in G(T), \forall A \in 2^{PH} : \int g d(f(A)) = \int \chi_{C(T)} g d(f(A)) \end{array} \right\} \end{aligned}$$

¹¹The juxtaposition $\chi_{F(T)} g$ means pointwise multiplication.

where

$$\begin{aligned} C(\mathcal{P} S) &= \{\mu : C_S \rightarrow 2^{PH}\} \\ C(S \rightarrow T) &= \{\mu : (C_S \rightarrow C_T) \rightarrow 2^{PH}\} \\ G(\mathcal{P} S) &= (\llbracket S \rrbracket \rightarrow 2^{PH}) \rightarrow \mathbb{R} \\ G(S \rightarrow T) &= ((\llbracket S \rrbracket \rightarrow \llbracket T \rrbracket) \rightarrow 2^{PH}) \rightarrow \mathbb{R} \end{aligned}$$

Similar to Section 5.3 the idea is to restrict the denotations to those semantics objects that are continuous (in the input set—not the environment), and for those partial maps that have a function domain, the domain only contains continuous functions. Note that we do not require that the partial maps *themselves* are continuous. Indeed, this is clearly not the case (e.g. in the case of *produce* V and $\lambda x:S.C$). The denotations for non-thunk values do not contain input sets or maps, hence C_S is simply $\llbracket S \rrbracket$. Thunked computations contain computations, thus $C_{\mathcal{T}T} = C_T$. For computations of type T , C_T says that the semantics itself must be continuous, and any element of a partial map's domain must be continuous.

We prove the following theorem, of which Theorem 5.5 is a direct consequence:

THEOREM 5.7. *For all values V and computations C :*

$$\text{if } \Gamma \vdash_v V : S \text{ and } \Gamma \vdash_c C : T, \text{ then } \llbracket V \rrbracket \rho \in C_S \text{ and } \llbracket C \rrbracket \rho \in C_T$$

where $\rho \in \llbracket \Gamma \rrbracket$ such that if $x : S_x \in \Gamma$, then $\rho(x) \in C_{S_x}$.

5.5 Conservativity

The next theorem relates the semantics to the original PNK semantics, showing that our semantics behaves identical to the original PNK semantics on probabilistic computations.

THEOREM 5.8 (CONSERVATIVITY). *Let C be a closed probabilistic computation, and let $\llbracket C \rrbracket_{PNK}$ be the denotation of C in the probabilistic PNK semantics [Smolka et al. 2017b], re-translated into ω QBS, then for all $A \in 2^{PH}$:*

$$\llbracket C \rrbracket_{PNK} A = \llbracket C \rrbracket () A \gg \lambda \mu \rightarrow \text{return } (\mu(()))$$

Summary. We have a *well-defined* denotational semantics for $\text{Pl}\omega\text{NK}$. This semantics is a *conservative* extension of PNK's semantics. However, it is not clear how to compute this semantics. Recall that application $C V$ applies every *unique* $f \in \text{dom}(\mu)$ *exactly* once, in parallel. Because parallel composition is not idempotent [Foster et al. 2016], duplicate applications are not innocent. The situation for sequencing is analogous. To the best of our knowledge, there is no decidable procedure for this uniqueness problem. We could follow Smolka et al. [2017a, 2019] and remove *dup*, making the state-space finite and discrete. Although a perfectly valid solution, we believe that this restricts the properties we can model in our language too much (e.g., modelling latency is not possible). Instead, we restrict the type of parallel composition to $\mathcal{P} \mathbf{1}$, to ensure that $\text{dom}(\mu)$ contains a single element. This enables the compilation of closed $\text{Pl}\omega\text{NK}$ terms into PNK. The program is then approximated with PNK's approximation procedure. We discuss this approach in detail in Section 6.

6 COMPILATION TO PNK

Let us make precise the restriction referred to in the previous section. We replace the judgements $\Gamma \vdash_v V : S$ and $\Gamma \vdash_c C : T$ with $\Gamma \Vdash_v V : S$ and $\Gamma \Vdash_c C : T$. The rules for these judgements are analogous

to the original rules, except that we replace

$$\frac{\Gamma \vdash_c C_1 : T \quad \Gamma \vdash_c C_2 : T}{\Gamma \vdash_c C_1 \& C_2 : T} \quad \text{with} \quad \frac{\Gamma \Vdash_c C_1 : \mathcal{P} \mathbf{1} \quad \Gamma \Vdash_c C_2 : \mathcal{P} \mathbf{1}}{\Gamma \Vdash_c C_1 \& C_2 : \mathcal{P} \mathbf{1}}$$

The new rule restricts the type of parallel composition to $\mathcal{P} \mathbf{1}$. In other words, the value of $\&$ is completely predictable: it must be *unit*. Furthermore, parallelism is the only way to grow the domain of the finite maps. Thus, all domains are now either a single function, or a single value. From a different perspective, we have just restricted the parallelism of Pl ω NK to the parallelism present in PNK.

6.1 Elaboration

We only compile closed computations of type $\mathcal{P} \mathbf{1}$. This is reasonable because complete Pl ω NK models are not functions and have no free variables. Initially, we only deal with *terminal* computations (Section 3.1). By case analysis on the typing judgement, such computations are either *produce unit* or probabilistic computations (i.e. only consist of PNK terms). The following relation elaborates these computations into PNK:

Definition 6.1 (Elaboration). Define $R \rightsquigarrow R$ as:

$$\begin{aligned} R &\rightsquigarrow R && \text{if } R \text{ is atomic} \\ R_1 ; R_2 &\rightsquigarrow E_1 ; E_2 && \text{if } R_1 \rightsquigarrow E_1 \text{ and } R_2 \rightsquigarrow E_2 \\ R_1 \& R_2 &\rightsquigarrow E_1 \& E_2 && \text{if } R_1 \rightsquigarrow E_1 \text{ and } R_2 \rightsquigarrow E_2 \\ R_1 \oplus R_2 &\rightsquigarrow E_1 \oplus E_2 && \text{if } R_1 \rightsquigarrow E_1 \text{ and } R_2 \rightsquigarrow E_2 \\ R^* &\rightsquigarrow E^* && \text{if } R \rightsquigarrow E \\ \text{produce } V &\rightsquigarrow \text{skip} \\ \lambda x : S. C &\rightsquigarrow \text{skip} \end{aligned}$$

Elaboration preserves the semantics of a closed term of type $\mathcal{P} \mathbf{1}$:

THEOREM 6.2 (SOUNDNESS OF ELABORATION). *Let R_1, R_2 be terminals such that $\Vdash_c R_1 : \mathcal{P} \mathbf{1}$, $\Vdash_c R_2 : \mathcal{P} \mathbf{1}$ and $R_1 \rightsquigarrow R_2$, then $\llbracket R_1 \rrbracket = \llbracket R_2 \rrbracket$.*

Moreover, the elaboration always exists for closed terms of the right type:

THEOREM 6.3 (COMPLETENESS OF ELABORATION). *Let R be a terminal, then there exists precisely one probabilistic E such that $R \rightsquigarrow E$.*

6.2 Reduction

Converting a closed terminal into a PNK program is only half the battle. The other half is performed by the bigstep relation $C \Downarrow R$ given in Figure 6. It reduces a computation to a terminal. We can make the following observations about the bigstep relation:

- Atomic computations simply reduce to themselves.
- Compound computations built with sequential and parallel composition, probabilistic choice and iteration are reduced to the reduction of their subcomputations.
- Forcing a thunk reduces to the reduct of the thunked computation.
- Application of a lambda substitutes the value into the body of the lambda. When the first argument reduces to sequential composition or probabilistic choice, the application *distributes* over this argument. This is *not* possible for parallel composition, since it is not distributive [Foster et al. 2016].
- When C_1 reduces to *produce* V , sequencing (C_1 to $x.C_2$) substitutes V for x in C_2 , reducing the result of the substitution. When the first argument instead reduces to a predicate, a

$$\boxed{C \Downarrow R}$$

$$\begin{array}{c}
\overline{P \Downarrow P} \quad \overline{F \leftarrow N \Downarrow F \leftarrow N} \quad \overline{dup \Downarrow dup} \quad \overline{produce V \Downarrow produce V} \quad \overline{\lambda x:S.C \Downarrow \lambda x:S.C} \\
\\
\frac{C_1 \Downarrow R_1 \quad C_2 \Downarrow R_2}{C_1 ; C_2 \Downarrow R_1 ; R_2} \quad \frac{C_1 \Downarrow R_1 \quad C_2 \Downarrow R_2}{C_1 \& C_2 \Downarrow R_1 \& R_2} \quad \frac{C_1 \Downarrow R_1 \quad C_2 \Downarrow R_2}{C_1 \oplus C_2 \Downarrow R_1 \oplus R_2} \quad \frac{C \Downarrow R}{C^* \Downarrow R^*} \\
\\
\frac{C \Downarrow R}{force(thunk C) \Downarrow R} \quad \frac{C_1 \Downarrow \lambda x:S.C_{11} \quad [x \mapsto V]C_{11} \Downarrow R}{C_1 V \Downarrow R} \quad \frac{C_1 \Downarrow R_{11} ; R_{12} \quad R_{12} V \Downarrow R_2}{C_1 V \Downarrow R_{11} ; R_2} \\
\\
\frac{C_1 \Downarrow produce V \quad [x \mapsto V]C_2 \Downarrow R}{C_1 to x.C_2 \Downarrow R} \quad \frac{C_1 \Downarrow R_{11} \oplus R_{12} \quad R_{11} V \Downarrow R_1 \quad R_{12} V \Downarrow R_2}{C_1 V \Downarrow R_1 \oplus R_2} \\
\\
\frac{C_1 \Downarrow P \quad [x \mapsto unit]C_2 \Downarrow R}{C_1 to x.C_2 \Downarrow P ; R} \quad \frac{C_1 \Downarrow F \leftarrow N \quad [x \mapsto unit]C_2 \Downarrow R}{C_1 to x.C_2 \Downarrow F \leftarrow N ; R} \\
\\
\frac{C_1 \Downarrow dup \quad [x \mapsto unit]C_2 \Downarrow R}{C_1 to x.C_2 \Downarrow dup ; R} \quad \frac{C_1 \Downarrow R_{11}^* \quad [x \mapsto unit]C_2 \Downarrow R_2}{C_1 to x.C_2 \Downarrow R_{11}^* ; R_2} \\
\\
\frac{C_1 \Downarrow R_{11} ; R_{12} \quad R_{12} to x.C_2 \Downarrow R_2}{C_1 to x.C_2 \Downarrow R_{11} ; R_2} \quad \frac{C_1 \Downarrow R_{11} \& R_{12} \quad [x \mapsto unit]C_2 \Downarrow R_2}{C_1 to x.C_2 \Downarrow (R_{11} \& R_{22}) ; R_2} \\
\\
\frac{C_1 \Downarrow R_{11} \oplus R_{12} \quad R_{11} to x.C_2 \Downarrow R_{21} \quad R_{12} to x.C_2 \Downarrow R_{22}}{C_1 to x.C_2 \Downarrow R_{21} \oplus R_{22}}
\end{array}$$

Fig. 6. Rules for reduction from $\text{Pl}\omega\text{NK}$ to PNK.

modification, a duplication, parallel composition or an iteration, C_1 always reduces to a terminal of type $\mathcal{P} \mathbf{1}$. Hence, we substitute *unit* for V in these cases. When the first argument reduces to sequential composition or probabilistic choice, sequencing, like application also *distributes* over this argument. Finally, if the first argument reduces to parallel composition, we know that the value it produces *must* be unit, and so we can always substitute *unit* for x .

To be clear, our compilation strategy is as follows: (1) *Reduce* a $\text{Pl}\omega\text{NK}$ computation to a terminal, and (2) *elaborate* the remaining terminal into PNK. To ensure that our compilation delivers correct results, it remains to show that step (1) terminates, and that this step is sound. Soundness means that the denotational semantics of the program is preserved:

THEOREM 6.4 (SOUNDNESS OF REDUCTION). *Let C, R be computations, if $\Gamma \Vdash_c C:T$ and $C \Downarrow R$ then $\llbracket C \rrbracket = \llbracket R \rrbracket$.*

Termination is our subsequent concern. It follows from strong normalisation of the reduction relation, a property of the meta-theory of reduction, discussed in the next section.

$$\mathbf{V}[\mathbf{1}] = \{\text{unit}\} \quad \mathbf{V}[\mathcal{H}] = \text{Headers} \quad \mathbf{V}[\mathbb{N}] = \mathbb{N} \quad \mathbf{V}[\mathcal{T}T] = \{\text{thunk } C \mid C \in \mathbf{C}[T]\}$$

$$\begin{aligned} R \in \mathbf{T}[\mathcal{P} \mathbf{1}] & \text{ iff } \Vdash_c R : \mathcal{P} \mathbf{1} \text{ where } R \text{ is atomic} \\ R^* \in \mathbf{T}[\mathcal{P} \mathbf{1}] & \text{ iff } R \in \mathbf{T}[\mathcal{P} \mathbf{1}] \text{ and } \Vdash_c R^* : \mathcal{P} \mathbf{1} \\ \text{produce } V \in \mathbf{T}[\mathcal{P} S] & \text{ iff } V \in \mathbf{V}[S] \text{ and } \Vdash_c \text{produce } V : \mathcal{P} S \\ R_1 ; R_2 \in \mathbf{T}[\mathcal{P} S] & \text{ iff } \exists T' : R_1 \in \mathbf{T}[T'] \text{ and } R_2 \in \mathbf{T}[\mathcal{P} S] \text{ and } \Vdash_c R_1 ; R_2 : \mathcal{P} S \\ R_1 \& R_2 \in \mathbf{T}[\mathcal{P} S] & \text{ iff } R_1, R_2 \in \mathbf{T}[\mathcal{P} S] \text{ and } \Vdash_c R_1 \& R_2 : \mathcal{P} S \\ R_1 \oplus R_2 \in \mathbf{T}[\mathcal{P} S] & \text{ iff } R_1, R_2 \in \mathbf{T}[\mathcal{P} S] \text{ and } \Vdash_c R_1 \oplus R_2 : \mathcal{P} S \\ R \in \mathbf{T}[S \rightarrow T] & \text{ iff } R \text{ is terminal and } \Vdash_c R : S \rightarrow T \text{ and } \forall V \in \mathbf{V}[S] : (R V) \in \mathbf{C}[T] \end{aligned}$$

$$C \in \mathbf{C}[T] \text{ iff } \Vdash_c C : T \text{ and } \exists R \in \mathbf{T}[T] : C \Downarrow R$$

Fig. 7. Logical relations involved in proving Strong Normalisation

6.3 Meta-theory of Reduction

The reduction relation obeys the standard type-preservation theorems, and in addition, is *strongly normalising*. In detail, terminals are reduced to themselves (Theorem 6.5). Moreover, all reductions result in terminals (by definition), in a deterministic fashion (Theorem 6.6) and preserve types (Theorem 6.7). Finally, reduction is strongly normalising (Theorem 6.8).

✓THEOREM 6.5 (REFLECTION). *Let R be a terminal, then $R \Downarrow R$.*

✓THEOREM 6.6 (DETERMINACY). *Let C, R_1, R_2 be computations, if $C \Downarrow R_1$ and $C \Downarrow R_2$, then $R_1 = R_2$.*

✓THEOREM 6.7 (PRESERVATION). *Let C, R be computations, if $\Gamma \Vdash_c C : T$ and $C \Downarrow R$, then $\Gamma \vdash_c R : T$. Conversely, if $\Gamma \Vdash_c C : T_1$, $\Gamma \Vdash_c R : T_2$ and $C \Downarrow R$, then $T_1 = T_2$.*

✓THEOREM 6.8 (STRONG NORMALISATION). *Let C be a computation such that $\Vdash_c C : T$ for some computation type T , then there exists a terminal R such that $C \Downarrow R$.*

In addition to the theorems shown here, Pl ω NK also satisfies additional inversion and substitution lemmas that are instrumental in proving these theorems.

The meta-theory discussed in this section has been mechanised with the aid of the Abella proof-assistant [Gacek 2008]. The proofs for these theorems proceed by induction either on the structure of terminals or the structure of the bigstep relation. Each of proof has many cases that need to be checked. By using a theorem prover, we ensure that no cases or conditions are forgotten.

The most involved proof is Strong Normalisation, which requires a logical-relation style proof technique [Tait 1967]. This particular proof was inspired by the standard proof of strong normalisation of CBPV, described in Levy's thesis [2001].

The definition of the logical relations are shown in Figure 7. In essence, we need to define three mutually-recursive type-indexed logical relations: one for values ($\mathbf{V}[S]$), one for closed *terminal* computations ($\mathbf{T}[T]$) and one for all computations, terminal or non-terminal ($\mathbf{C}[T]$). The idea is that $\mathbf{C}[T]$ contains only closed computations for which the bigstep relation $\Downarrow$ terminates, and only produces terminating values. For values, $\mathbf{V}[S]$ contains all closed non-thunk values, and only closed thunks of terminating computations. Additionally, computations of function type must preserve termination when applied to terminating values.

Compared to Levy's logical relations, our logical relation does not contain product or sum types. In our proof we leverage standard Abella techniques to prove theorems with logical relations, which requires defining and proving substitution lemmas for every syntactic form. Moreover, we also need to show additional preservation lemmas for sequential and parallel composition, and for probabilistic choice.

The proof of strong normalisation, together with the definitions of the logical relations and supporting lemmas amounts to a little under 800 lines of Abella code (roughly 1500 LOC in total).

6.4 Discussion

We have identified a class of $\text{Pl}\omega\text{NK}$ programs that can be safely compiled to PNK. In particular, the class consists of computations C such that $\Vdash_c C : \mathcal{P} \mathbf{1}$, i.e., those programs that do not use parallelism beyond what is present in PNK. It should be possible to ease this restriction slightly. For instance, unlike functions, we can distinguish distinct headers and literals. This leads to only finitely many cases, which could be encoded explicitly into PNK.

Once a program has been compiled, it can be approximated in PNK. The approximation proceeds by expanding iterations in the program up to n times, for finite n . More iterations improve the accuracy of the results [Smolka et al. 2017b].

7 RELATED WORK

Software Defined Networks. Software Defined Networks (SDN) [Foster et al. 2013] aim to decrease the complexity of the modern computer networking environment, by offering a clean open interface between heterogeneous networking devices (e.g. routers, switches and firewalls). This is accomplished through the OpenFlow¹² protocol. Unfortunately, this is a rather low-level protocol making it inconvenient to program hardware in OpenFlow directly.

The aim of the Frenetic project [Reich et al. 2013] is to design the right high-level abstractions for controlling OpenFlow hardware. NetKAT [Anderson et al. 2014] and PNK [Foster et al. 2016] were developed under this project. Some preliminary case studies were performed with PNK, modelling the behaviour of several traffic engineering approaches. The effort involved in these case studies was not reported. Interestingly, the few samples of code that were provided seem to use features (e.g. finite iterations, variables), that are not part of the formalised fragment of PNK, but could be implemented fully within $\text{Pl}\omega\text{NK}$.

Probabilistic Programming. The goal of probabilistic programming [Goodman 2013] can be captured by the following equation:

$$\text{PPL} = \text{MODELLING LANGUAGE} + \text{INFERENCE ALGORITHM}$$

That is, probabilistic programming's goal is to unify probabilistic modelling and general purpose programming: probabilistic models are written in the language, and the probabilities are inferred using a generic inference algorithm.

Stan [Carpenter et al. 2017] is a very popular statistical modelling language, with bindings to R, Python, MATLAB, Julia and several others. More recently, languages such as Gen [Cusumano-Towner et al. 2019] and Turing [Ge et al. 2018] have started to make the inference algorithms programmable, in addition to the model, since fine-tuning the inference can lead to large performance gains.

Probabilistic Programming has been applied to such diverse problems as 3D body pose estimation from depth data [Cusumano-Towner et al. 2019], genetics [De Maeyer et al. 2013] and Automatic Video Montage [Aerts et al. 2016].

¹²<https://www.opennetworking.org/>

Functional PPLs such as Anglican [Wood et al. 2014], Venture [Lu 2016], or Gen [Lu 2016], allow distributions over higher-order functions. However, unlike this work, they do not formalise the semantics of higher-order functions, focusing instead on language design and implementation, and inference. The work on quasi-Borel Spaces is at least partially motivated by the unfilled need for a theoretical foundation for these languages [Heunen et al. 2017]. Quasi-Borel Spaces have been used by Scibior et al. [2018] to verify the correctness of modular Bayesian inference algorithms.

Probabilistic powerdomains have been extensively investigated (see e.g., Bacci et al. [2018]; Battenfeld et al. [2007]; Goubault-Larrecq and Varacca [2011]; Jones and Plotkin [1989]; Jung and Tix [1998]; Saheb-Djahromi [1980]). Nevertheless, until the work of Vákár et al. [2019], a convenient continuous probabilistic powerdomain that supports iteration and is commutative proved elusive. However, the ω -quasi Borel Spaces are not the only approach to this problem, as we remark in the the next paragraph.

Probabilistic Call-By-Push-Value. Although $\text{Pl}\omega\text{NK}$ does not model CBPV exactly, it was heavily inspired by it. CBPV was developed by Levy [2001] as a paradigm that subsumes both CBV and CBN. Since then, Ehrhard and Tasson [2019] have developed a probabilistic CBPV calculus. Technically, they give a semantics in terms of probabilistic coherence spaces [Danos and Ehrhard 2011], whereas we use a monadic semantics based on ωQBSes .

Goubault-Larrecq [2019] cleverly side-steps the issue of providing a commutative statistical higher-order powerdomain, by giving a semantics for CBPV that interprets value types and computation types differently. The values are interpreted in a category which is closed under the powerdomain functor, and the computations are interpreted as DCPOs. His language also has demonic non-determinism, statistical termination testers and parallel if statements.

A crucial difference with our work is that monadic state is not present in either calculus, while it is in ours. This significantly complicates our denotational semantics, in particular for application and sequencing. Moreover, these calculi do observe the isomorphism mentioned in Section 5.1. A detailed investigation into the relationships between these calculi and our language is reserved for future work.

8 CONCLUSIONS AND FUTURE WORK

In future work, we intend to quantitatively evaluate the impact $\text{Pl}\omega\text{NK}$, by re-implementing the existing PNK case-studies, and study the improvements in terms of readability and maintainability. To support this quantitative study, we need to further develop and optimise our prototype. We believe that performance can be improved by incorporating techniques such as knowledge compilation [Kisa et al. 2014; Smolka et al. 2015] and defunctionalisation [Danvy and Nielsen 2001; Reynolds 1998].

Furthermore, our work contains a significant amount of paper proofs. These proofs are inductive proofs, using equational reasoning and logical relations, which should not be too difficult to mechanise. However, the requisite background theory, i.e., $(\omega\text{-})$ Quasi-Borel Spaces, has to be mechanised first.

Lastly, we are investigating reformulations of $\text{Pl}\omega\text{NK}$ for other paradigms such as true CBPV and Fine-Grained CBV [Levy 2001, App. A.3]. The key difference appears to be that function typed terms should only evaluate their side-effects when applied to a value.

Conclusion. In this article we presented $\text{Pl}\omega\text{NK}$, a functional network modelling language that combines state, parallelism and probabilistic choice. Because it combines higher-order functions and probability, we cannot give it a purely measure-theoretic semantics. Instead, we leverage ω -Quasi Borel Spaces to define our denotational semantics. We also define a strongly normalising type system for $\text{Pl}\omega\text{NK}$. Since the main purpose of $\text{Pl}\omega\text{NK}$ is verification, the additional flexibility of

general recursion is not required. Indeed, strong normalisation is necessary to make our semantics well-defined. Moreover, we develop a procedure to compile programs in our language to the simpler language Probabilistic NetKAT, given small type restrictions.

ACKNOWLEDGMENTS

We would like to thank Ohad Kammar and Mathijs Vákár for personally explaining ω QBSes and providing us with an early draft of their paper [Vákár et al. 2019].

We are also grateful to the anonymous POPL reviewers for their constructive feedback and helpful in-depth comments.

Alexander Vandenbroucke is an SB Fellow of the Flemish Fund for Scientific Research (FWO), File No.: 1S68117N. This work is further supported by FWO Grant No. G095917N.

REFERENCES

- Bram Aerts, Toon Goedemé, and Joost Vennekens. 2016. A Probabilistic Logic Programming Approach to Automatic Video Montage. In *Proceedings of the Twenty-second European Conference on Artificial Intelligence (ECAI'16)*. IOS Press, Amsterdam, The Netherlands, The Netherlands, 234–242. <https://doi.org/10.3233/978-1-61499-672-9-234>
- Carolyn Jane Anderson, Nate Foster, Arjun Guha, Jean-Baptiste Jeannin, Dexter Kozen, Cole Schlesinger, and David Walker. 2014. NetKAT: semantic foundations for networks. In *POPL*. ACM, 113–126.
- Robert J Aumann et al. 1961. Borel structures for function spaces. *Illinois Journal of Mathematics* 5, 4 (1961), 614–630.
- Giorgio Bacci, Robert Furber, Dexter Kozen, Radu Mardare, Prakash Panangaden, and Dana Scott. 2018. Boolean-Valued Semantics for the Stochastic λ -Calculus. In *LICS*. ACM, 669–678.
- Ingo Battenfeld, Matthias Schröder, and Alex Simpson. 2007. A Convenient Category of Domains. *Electr. Notes Theor. Comput. Sci.* 172 (2007), 69–99.
- Bob Carpenter, Andrew Gelman, Matthew D Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. 2017. Stan: A probabilistic programming language. *Journal of statistical software* 76, 1 (2017).
- Marco F. Cusumano-Towner, Feras A. Saad, Alexander K. Lew, and Vikash K. Mansinghka. 2019. Gen: a general-purpose probabilistic programming system with programmable inference. In *PLDI*. ACM, 221–236.
- Vincent Danos and Thomas Ehrhard. 2011. Probabilistic coherence spaces as a model of higher-order probabilistic computation. *Inf. Comput.* 209, 6 (2011), 966–991.
- Olivier Danvy and Lasse R. Nielsen. 2001. Defunctionalization at Work. In *PPDP*. ACM, 162–174.
- Dries De Maeyer, Joris Renkens, Lore Cloots, Luc De Raedt, and Kathleen Marchal. 2013. PheNetic: network-based interpretation of unstructured gene lists in *E. coli*. *Molecular BioSystems* 9, 7 (2013), 1594–1603.
- Thomas Ehrhard and Christine Tasson. 2019. Probabilistic call by push value. *Logical Methods in Computer Science* 15, 1 (2019).
- Daan Fierens, Guy Van den Broeck, Joris Renkens, Dimitar Sht. Shterionov, Bernd Gutmann, Ingo Thon, Gerda Janssens, and Luc De Raedt. 2015. Inference and learning in probabilistic logic programs using weighted Boolean formulas. *TPLP* 15, 3 (2015), 358–401.
- Nate Foster, Arjun Guha, Mark Reitblatt, Alec Story, Michael J Freedman, Naga Praveen Katta, Christopher Monsanto, Joshua Reich, Jennifer Rexford, Cole Schlesinger, et al. 2013. Languages for software-defined networks. *IEEE Communications Magazine* 51, 2 (2013), 128–134.
- Nate Foster, Dexter Kozen, Konstantinos Mamouras, Mark Reitblatt, and Alexandra Silva. 2016. Probabilistic NetKAT. In *Programming Languages and Systems - 25th European Symposium on Programming, ESOP 2016, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2016, Eindhoven, The Netherlands, April 2-8, 2016, Proceedings (Lecture Notes in Computer Science)*, Peter Thiemann (Ed.), Vol. 9632. Springer, 282–309. https://doi.org/10.1007/978-3-662-49498-1_12
- Andrew Gacek. 2008. The Abella Interactive Theorem Prover (System Description). In *IJCAR (Lecture Notes in Computer Science)*, Vol. 5195. Springer, 154–161.
- Hong Ge, Kai Xu, and Zoubin Ghahramani. 2018. Turing: A language for flexible probabilistic inference. In *International Conference on Artificial Intelligence and Statistics*. 1682–1690.
- Noah D. Goodman. 2013. The principles and practice of probabilistic programming. In *POPL*. ACM, 399–402.
- Noah D. Goodman, Vikash K. Mansinghka, Daniel M. Roy, Keith Bonawitz, and Joshua B. Tenenbaum. 2012. Church: a language for generative models. *CoRR* abs/1206.3255 (2012).
- Jean Goubault-Larrecq. 2019. A Probabilistic and Non-Deterministic Call-by-Push-Value Language. In *Thirty-Fourth Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2019, Vancouver, Canada, June 24-27, 2019, Proceedings*. To appear.

<https://arxiv.org/pdf/1812.11573.pdf>

- Jean Goubault-Larrecq and Daniele Varacca. 2011. Continuous Random Variables. In *LICS*. IEEE Computer Society, 97–106.
- Chris Heunen, Ohad Kammar, Sam Staton, and Hongseok Yang. 2017. A convenient category for higher-order probability theory. In *LICS*. IEEE Computer Society, 1–12.
- Claire Jones and Gordon D Plotkin. 1989. A probabilistic powerdomain of evaluations. In *[1989] Proceedings. Fourth Annual Symposium on Logic in Computer Science*. IEEE, 186–195.
- Achim Jung and Regina Tix. 1998. The troublesome probabilistic powerdomain. *Electr. Notes Theor. Comput. Sci.* 13 (1998), 70–91.
- Doga Kisa, Guy Van den Broeck, Arthur Choi, and Adnan Darwiche. 2014. Probabilistic Sentential Decision Diagrams. In *KR*. AAAI Press.
- Dexter Kozen. 1981. Semantics of Probabilistic Programs. *J. Comput. Syst. Sci.* 22, 3 (1981), 328–350.
- Paul Blain Levy. 2001. *Call-by-push-value*. Ph.D. Dissertation. Queen Mary University of London, UK.
- Anthony Lu. 2016. *Venture: An extensible platform for probabilistic meta-programming*. Master’s thesis. Massachusetts Institute of Technology.
- Eugenio Moggi. 1991. Notions of Computation and Monads. *Inf. Comput.* 93, 1 (1991), 55–92.
- Benjamin C. Pierce. 2002. *Types and programming languages*. MIT Press.
- Joshua Reich, Christopher Monsanto, Nate Foster, Jennifer Rexford, and David Walker. 2013. Modular SDN programming with Pyretic. *Technical Report of USENIX* (2013).
- John C. Reynolds. 1998. Definitional Interpreters for Higher-Order Programming Languages. *Higher-Order and Symbolic Computation* 11, 4 (1998), 363–397.
- Nasser Saheb-Djahromi. 1980. CPO’s of measures for nondeterminism. *Theoretical Computer Science* 12, 1 (1980), 19–37.
- Adam Scibior, Ohad Kammar, and Zoubin Ghahramani. 2018. Functional programming for modular Bayesian inference. *PACMPL* 2, ICFP (2018), 83:1–83:29.
- Steffen Smolka, Spiridon Aristides Eliopoulos, Nate Foster, and Arjun Guha. 2015. A fast compiler for NetKAT. In *ICFP*. ACM, 328–341.
- Steffen Smolka, David M. Kahn, Praveen Kumar, Nate Foster, Dexter Kozen, and Alexandra Silva. 2017a. Deciding Probabilistic Program Equivalence in NetKAT. *CoRR* abs/1707.02772 (2017). arXiv:1707.02772 <http://arxiv.org/abs/1707.02772>
- Steffen Smolka, Praveen Kumar, Nate Foster, Dexter Kozen, and Alexandra Silva. 2017b. Cantor meets Scott: semantic foundations for probabilistic networks. In *POPL*. ACM, 557–571.
- Steffen Smolka, Praveen Kumar, David M. Kahn, Nate Foster, Justin Hsu, Dexter Kozen, and Alexandra Silva. 2019. Scalable verification of probabilistic networks. In *PLDI*. ACM, 190–203.
- Sam Staton. 2017. Commutative Semantics for Probabilistic Programming. In *ESOP (Lecture Notes in Computer Science)*, Vol. 10201. Springer, 855–879.
- William W. Tait. 1967. Intensional Interpretations of Functionals of Finite Type I. *J. Symb. Log.* 32, 2 (1967), 198–212.
- Matthijs Vákár, Ohad Kammar, and Sam Staton. 2019. A domain theory for statistical probabilistic programming. *PACMPL* 3, POPL (2019), 36:1–36:29. <https://dl.acm.org/citation.cfm?id=3290349>
- Frank D. Wood, Jan-Willem van de Meent, and Vikash Mansinghka. 2014. A New Approach to Probabilistic Programming Inference. In *AISTATS (JMLR Workshop and Conference Proceedings)*, Vol. 33. JMLR.org, 1024–1032.