Etalumis: Bringing Probabilistic Programming to Scientific Simulators at Scale

Atılım Güneş Baydin
University of Oxford

Lei Shao
Intel Corporation

Wahid Bhimji
Lawrence Berkeley National
Laboratory

Lukas Heinrich
CERN

Lawrence Meadows
Intel Corporation

Jialin Liu
Lawrence Berkeley National
Laboratory

Andreas Munk
University of British Columbia

Saeid Naderiparizi
University of British Columbia

Bradley Gram-Hansen
University of Oxford

Gilles Louppe
University of Liège

Mingfei Ma
Intel Corporation

Xiaohui Zhao
Intel Corporation

Philip Torr
University of Oxford

Victor Lee
Intel Corporation

Kyle Cranmer
New York University

Prabhat
Lawrence Berkeley National
Laboratory

Frank Wood
University of British Columbia

ABSTRACT

Probabilistic programming languages (PPLs) are receiving widespread attention for performing Bayesian inference in complex generative models. However, applications to science remain limited because of the impracticability of rewriting complex scientific simulators in a PPL, the computational cost of inference, and the lack of scalable implementations. To address these, we present a novel PPL framework that couples directly to existing scientific simulators through a cross-platform probabilistic execution protocol and provides Markov chain Monte Carlo (MCMC) and deep-learning-based inference compilation (IC) engines for tractable inference. To guide IC inference, we perform distributed training of a dynamic 3DCNN-LSTM architecture with a PyTorch-MPI-based framework on 1,024 32-core CPU nodes of the Cori supercomputer with a global mini-batch size of 128k: achieving a performance of 450 Tflop/s through enhancements to PyTorch. We demonstrate a Large Hadron Collider (LHC) use-case with the C++ Sherpa simulator and achieve the largest-scale posterior inference in a Turing-complete PPL.

KEYWORDS

probabilistic programming, simulation, inference, deep learning

ACM Reference Format:

Atılım Güneş Baydin, Lei Shao, Wahid Bhimji, Lukas Heinrich, Lawrence Meadows, Jialin Liu, Andreas Munk, Saeid Naderiparizi, Bradley Gram-Hansen, Gilles Louppe, Mingfei Ma, Xiaohui Zhao, Philip Torr, Victor Lee, Kyle Cranmer, Prabhat, and Frank Wood. 2019. Etalumis: Bringing Probabilistic Programming to Scientific Simulators at Scale. In *The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC '19)*, November 17–22, 2019, Denver, CO, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3295500.3356180>

1 INTRODUCTION

Probabilistic programming [71] is an emerging paradigm within machine learning that uses general-purpose programming languages to express probabilistic models. This is achieved by introducing statistical conditioning as a language construct so that inverse problems can be expressed. Probabilistic programming languages (PPLs) have semantics [67] that can be understood as Bayesian inference [13, 24, 26]. The major challenge in designing useful PPL systems is that language evaluators must solve arbitrary, user-provided inverse problems, which usually requires general-purpose inference algorithms that are computationally expensive.

In this paper we report our work that enables, for the first time, the use of *existing* stochastic simulator code as a probabilistic program in which one can do fast, repeated (amortized) Bayesian inference; this enables one to predict the distribution of input parameters and all random choices in the simulator from an observation of its output. In other words, given a simulator of a generative process in the forward direction (inputs→outputs), our technique can provide the reverse (outputs→inputs) by predicting the whole latent state of the simulator that could have given rise to an observed instance of its output. For example, using a particle physics simulation we

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
SC '19, November 17–22, 2019, Denver, CO, USA
© 2019 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-6229-0/19/11.
<https://doi.org/10.1145/3295500.3356180>

can get distributions over the particle properties and decays within the simulator that can give rise to a collision event observed in a detector, or, using a spectroscopy simulator we can determine the elemental matter composition and dispersions within the simulator explaining an observed spectrum. In fields where accurate simulators of real-world phenomena exist, our technique enables the interpretable explanation of real observations under the structured model defined by the simulator code base.

We achieve this by defining a probabilistic programming execution protocol that interfaces with existing simulators at the sites of random number draws, without altering the simulator’s structure and execution in the host system. The random number draws are routed through the protocol to a PPL system which treats these as samples from corresponding prior distributions in a Bayesian setting, giving one the capability to record or guide the execution of the simulator to perform inference. Thus we generalize existing simulators as probabilistic programs and make them subject to inference under general-purpose inference engines.

Inference in the probabilistic programming setting is performed by sampling in the space of execution traces, where a single sample (an execution trace) represents a full run of the simulator. Each execution trace itself is composed of a potentially unbounded sequence of addresses, prior distributions, and sampled values, where an address is a unique label identifying each random number draw. In other words, we work with empirical distributions over simulator executions, which entails unique requirements on memory, storage, and computation that we address in our implementation. The addresses comprising each trace give our technique the unique ability to provide direct connections to the simulator code base for any predictions at test time, where the simulator is no longer used as a black box but as a highly structured and interpretable probabilistic generative model that it implicitly represents.

Our PPL provides inference engines from the Markov chain Monte Carlo (MCMC) and importance sampling (IS) families. MCMC inference guarantees closely approximating the true posterior of the simulator, albeit with significant computational cost due to its sequential nature and the large number of iterations one needs to accumulate statistically independent samples. Inference compilation (IC) [47] addresses this by training a dynamic neural network to provide proposals for IS, leading to fast amortized inference.

We name this project “Etalumis”, the word “simulate” spelled backwards, as a reference to the fact that our technique essentially inverts a simulator by probabilistically inferring all choices in the simulator given an observation of its output. We demonstrate this by inferring properties of particles produced at the Large Hadron Collider (LHC) using the Sherpa¹ [29] simulator.

1.1 Contributions

Our main contributions are:

- A novel PPL framework that enables execution of existing stochastic simulators under the control of general-purpose inference engines, with HPC features including handling multi-TB data and distributed training and inference.
- The largest scale posterior inference in a Turing-complete PPL, where our experiments encountered approximately 25,000 latent

variables² expressed by the existing Sherpa simulator code base of nearly one million lines of code in C++ [29].

- Synchronous data parallel training of a dynamic 3DCNN–LSTM neural network (NN) architecture using the PyTorch [61] MPI framework at the scale of 1,024 nodes (32,768 CPU cores) with a global minibatch size of 128k. To our knowledge this is the largest scale use of PyTorch’s builtin MPI functionality,³ and the largest minibatch size used for this form of NN model.

2 PROBABILISTIC PROGRAMMING FOR PARTICLE PHYSICS

Particle physics seeks to understand particles produced in collisions at accelerators such as the LHC at CERN. Collisions happen millions of times per second, creating cascading particle decays, observed in complex instruments such as the ATLAS detector [2], comprising millions of electronics channels. These experiments analyze the vast volume of resulting data and seek to reconstruct the initial particles produced in order to make discoveries including physics beyond the current Standard Model of particle physics [28][73][63][72].

The Standard Model has a number of parameters (e.g., particle masses), which we can denote θ , describing the way particles and fundamental forces act in the universe. In a given collision at the LHC, with initial conditions denoted E , we observe a cascade of particles interact with particle detectors. If we denote *all* of the random “choices” made by nature as $\mathbf{x}$, the Standard Model describes, generatively, the conditional probability $p(\mathbf{x}|E, \theta)$, that is, the distribution of all choices $\mathbf{x}$ as a function of initial conditions E and model parameters θ . Note that, while the Standard Model can be expressed symbolically in mathematical notation [32, 62], it can also be expressed computationally as a stochastic simulator [29], which, given access to a random number generator, can draw samples from $p(\mathbf{x})$.⁴ Similarly, a particle detector can be modeled as a stochastic simulator, generating samples from $p(\mathbf{y}|\mathbf{x})$, the likelihood of observation $\mathbf{y}$ as a function of $\mathbf{x}$.

In this paper we focus on a real use-case in particle physics, performing experiments on the decay of the τ (tau) lepton. This is under active investigation by LHC physicists [4] and important to uncovering properties of the Higgs boson. We use the state-of-the-art Sherpa simulator [29] for modeling τ particle creation in LHC collisions and their subsequent decay into further particles (the stochastic events $\mathbf{x}$ above), coupled to a fast 3D detector simulator for the detector observation $\mathbf{y}$.

Current methods in the field include performing classification and regression using machine learning approaches on low dimensional distributions of derived variables [4] that provide point-estimates without the posterior of the full latent state nor the deep interpretability of our approach. Inference of the latent structure has only previously been used in the field with drastically simplified models of the process and detector [43] [3].

PPLs allow us to express inference problems such as: given an actual particle detector observation $\mathbf{y}$, what sequence of choices $\mathbf{x}$ are likely to have led to this observation? In other words, we would

²Note that the simulator defines an unlimited number of random variables because of the presence of rejection sampling loops.

³Personal communication with PyTorch developers.

⁴Dropping the dependence on E and θ because everything in this example is conditionally dependent on these quantities.

¹<https://gitlab.com/sherpa-team/sherpa>

like to find $p(\mathbf{x}|\mathbf{y})$, the distribution of $\mathbf{x}$ as a function of $\mathbf{y}$. To solve this inverse problem via conditioning requires invoking Bayes rule

$$p(\mathbf{x}|\mathbf{y}) = \frac{p(\mathbf{y}, \mathbf{x})}{p(\mathbf{y})} = \frac{p(\mathbf{y}|\mathbf{x})p(\mathbf{x})}{\int p(\mathbf{y}|\mathbf{x})p(\mathbf{x})d\mathbf{x}}$$

where the posterior distribution of interest, $p(\mathbf{x}|\mathbf{y})$, is related to the composition of the two stochastic simulators in the form of the joint distribution $p(\mathbf{y}, \mathbf{x}) = p(\mathbf{y}|\mathbf{x})p(\mathbf{x})$ renormalized by the marginal probability, or evidence of the data, $p(\mathbf{y}) = \int p(\mathbf{y}|\mathbf{x})p(\mathbf{x})d\mathbf{x}$. Computing the evidence requires summing over all possible paths that the simulation can take. This is a large number of possible paths; in most models this is a quantity that is impossible to compute in polynomial time. In practice PPLs approximate the posterior $p(\mathbf{x}|\mathbf{y})$ using sampling-based inference engines that sidestep the integration problem but remain computationally intensive. This specifically is where probabilistic programming meets, for the first time in this paper, high-performance computing.

3 STATE OF THE ART

3.1 Probabilistic programming

Within probabilistic programming, recent advances in computational hardware have made it possible to distribute certain types of inference processes, enabling inference to be applied to problems of real-world relevance [70]. By parallelizing computation over several cores, PPLs have been able to perform large-scale inference on models with increasing numbers of observations, such as the cause and effect analysis of 1.6×10^9 genetic measurements [30, 70], spatial analysis of 1.5×10^4 shots from 308 NBA players [18], exploratory analysis of 1.7×10^6 taxi trajectories [36], and probabilistic modeling for processing hundreds-of-thousands of Xbox live games per day to rank and match players fairly [33, 55].

In all these large-scale programs, despite the number of observations being large, model sizes in terms of the number of latent variables have been limited [36]. In contrast, to perform inference in a complex scientific model such as the Standard Model encoded by Sherpa requires handling thousands of latent variables, all of which need to be controlled within the program to perform inference in a scalable manner. To our knowledge, no existing PPL system has been used to run inference at the scale we are reporting in this work, and instances of distributed inference in existing literature have been typically restricted to small clusters [19].

A key feature of PPLs is that they decouple model specification from inference. A model is implemented by the user as a stand-alone regular program in the host programming language, specifying a generative process that produces samples from the joint prior distribution $p(\mathbf{y}, \mathbf{x}) = p(\mathbf{y}|\mathbf{x})p(\mathbf{x})$ in each execution, that is, a forward model going from choices $\mathbf{x}$ to outcomes (observations) $\mathbf{y}$. The same program can then be executed using a variety of general-purpose inference engines available in the PPL system to obtain $p(\mathbf{x}|\mathbf{y})$, the inverse going from observations $\mathbf{y}$ to choices $\mathbf{x}$. Inference engines available in PPLs range from MCMC-based lightweight Metropolis Hastings (LMH) [74] and random-walk Metropolis Hastings (RMH) [46] algorithms to importance sampling (IS) [8] and sequential Monte Carlo [22]. Modern PPLs such as Pyro [11] and TensorFlow Probability [19, 70] use gradient-based inference engines including variational inference [36, 42] and Hamiltonian Monte Carlo [37, 57]

that benefit from modern deep learning hardware and automatic differentiation [9] features provided by PyTorch [61] and TensorFlow [5] libraries. Another way of making use of gradient-based optimization is to combine IS with deep-learning-based proposals trained with data sampled from the probabilistic program, resulting in the IC algorithm [47, 49] in an amortized inference setting [25].

3.2 Distributed training for deep learning

To perform IC inference in Turing-complete PPLs in general, we would like to support the training of dynamic NNs whose runtime structure changes in each execution of the probabilistic model by rearranging NN modules corresponding to different addresses (unique random number draws) encountered [47] (Section 4.3). Moreover, depending on probabilistic model complexity, the NNs may grow in size if trained in an online setting, as a model can represent a potentially unbounded number of random number draws. In addition to these, the volume of training data required is large, as the data keeps track of all execution paths within the simulator. To enable rapid prototyping, model evaluation, and making use of HPC capacity, scaling deep learning training to multiple computation units is highly desirable [38, 44, 45, 51, 52].

In this context there are three prominent parallelism strategies: data- and model-parallelism, and layer pipelining. In this project we work in a data-parallel setting where different nodes train the same model on different subsets of data. For such training, there are synchronous- and asynchronous-update approaches. In synchronous update [16, 59], locally computed gradients are summed across the nodes at the same time with synchronization barriers for parameter update. In asynchronous update [17, 58, 77], one removes the barrier so that nodes can independently contribute to a parameter server. Although synchronous update can entail challenges due to straggler effects [15, 69], it has desirable properties in terms of convergence, reproducibility, and ease of debugging. In this work, given the novelty of the probabilistic techniques we are introducing and the need to fully understand and compare trained NNs without ambiguity, we employ synchronous updates.

In synchronous updates, large global minibatches can make convergence challenging and hinder test accuracy. Keskar et al. [39] pointed out large-minibatch training can lead to sharp minima and a generalization gap. Other work [31, 76] argues that the difficulties in large-minibatch training are optimization related and can be mitigated with learning rate scaling [31]. You et al. [76] apply layer-wise adaptive rate scaling (LARS) to achieve large-minibatch-size training of a Resnet-50 architecture without loss of accuracy, and Ginsburg et al. [27] use layer-wise adaptive rate control (LARC) to improve training stability and speed. Smith et al. [65] have proposed to increase the minibatch size instead of decaying the learning rate, and more recent work [53, 64] showed relationships between gradient noise scale (or training steps) and minibatch size. Through such methods, distributed training has been scaled to many thousands of CPUs or GPUs [44, 45, 51, 54]. While we take inspiration from these recent approaches, our dynamic NN architecture and training data create a distinct training setting which requires appropriate innovations, as discussed in Section 4.3.

Figure 1: The probabilistic execution protocol (PPX). *Sample* and *observe* statements correspond to random number draws and conditioning, respectively.

4 INNOVATIONS

4.1 PPX and pyprob: executing existing simulators as probabilistic programs

One of our main contributions in Etalumis is the development of a probabilistic programming execution protocol (PPX), which defines a cross-platform API for the execution and control of stochastic simulators⁵ (Figure 1). The protocol provides language-agnostic definitions of common probability distributions and message pairs covering the call and return values of: (1) program entry points; (2) *sample* statements for random number draws; and (3) *observe* statements for conditioning. The purpose of this protocol is twofold:

- It allows us to record execution traces of a stochastic simulator as a sequence of *sample* and *observe* (conditioning) operations on random numbers, each associated with an address A_t . We can use these traces for tasks such as inspecting the probabilistic model implemented by the simulator, computing likelihoods, learning surrogate models, and generating training data for IC NNs.
- It allows us to control the execution of the simulator, at inference time, by making intelligent choices for each random number draw as the simulator keeps requesting random numbers. General-purpose PPL inference guides the simulator by making random number draws not from the prior $p(\mathbf{x})$ but from proposal distributions $q(\mathbf{x}|\mathbf{y})$ that depend on observed data $\mathbf{y}$ (Section 2).

PPX is based on flatbuffers,⁶ a streamlined version of Google protocol buffers, providing bindings into C++, C#, Go, Java, JavaScript, PHP, Python, and TypeScript, enabling lightweight PPL front ends in these languages—in the sense of requiring the implementation of a simple intermediate layer to perform *sample* and *observe* operations over the protocol. We exchange PPX messages over ZeroMQ⁷ [34] sockets, which allow communication between separate processes in the same machine (via inter-process sockets) or across a network (via TCP). PPX is inspired by the Open Neural Network Exchange (ONNX) project⁸ allowing interoperability between major deep learning frameworks, and it allows the execution of any

stochastic simulator under the control of any PPL system, provided that the necessary bindings are incorporated on both sides.

Using the PPX protocol as the interface, we implement two main components: (1) pyprob, a PyTorch-based PPL⁹ in Python and (2) a C++ binding to the protocol to route the random number draws in Sherpa to the PPL and therefore allow probabilistic inference in this simulator. Our PPL is designed to work with models written in Python and other languages supported through PPX. This is in contrast to existing PPLs such as Pyro [11] and TensorFlow Probability [19, 70] which do not provide a way to interface with existing simulators and require one to implement any model from scratch in the specific PPL.¹⁰ We develop pyprob based on PyTorch [61], to utilize its automatic differentiation [9] infrastructure with support for dynamic computation graphs for IC inference.

4.2 Efficient Bayesian inference

Working with existing simulators as probabilistic programs restricts the class of inference engines that we can put to use. Modern PPLs commonly use gradient-based inference such as Hamiltonian Monte Carlo [57] and variational inference [36, 42] to approximate posterior distributions. However this is not applicable in our setting due to the absence of derivatives in general simulator codes. Therefore in pyprob we focus our attention on two inference engine families that can control Turing-complete simulators over the PPX protocol: MCMC in the RMH variety [46, 74], which provides a high-compute-cost sequential algorithm with statistical guarantees to closely approximate the posterior, and IS with IC [47], which does not require derivatives of the simulator code but still benefits from gradient-based methods by training proposal NNs and using these to significantly speed up IS inference.

It is important to note that the inference engines in pyprob work in the space of execution traces of probabilistic programs, such that a single sample from the inference engine corresponds to a full run of the simulator. Inference in this setting amounts to making adjustments to the random number draws, re-executing the simulator, and scoring the resulting execution in terms of the likelihood of the given observation. Depending on the specific observation and the simulator code involved, inference is computationally very expensive, requiring up to millions of executions in the RMH engine. Despite being very costly, RMH provides a way of sampling from the true posterior [56, 57], which is needed in initial explorations of any new simulator to establish correct posteriors serving as reference to confirm that IC inference can work correctly in the given setting. To establish the correctness of our inference results, we implement several MCMC convergence diagnostics. Autocorrelation measures the number of iterations one needs to get effectively independent samples in the same MCMC chain, which allows us to estimate how long RMH needs to run to reach a target effective sample size. The Gelman–Rubin metric, given multiple independent MCMC chains sampled from the same posterior, compares the variance of each chain to the pooled variance of all chains to statistically establish that we converged on the true posterior [24].

RMH comes with a high computational cost. This is because it requires a large number of initial samples to be generated that are

⁵<https://github.com/probprog/ppx>

⁶<http://google.github.io/flatbuffers/>

⁷<http://zeromq.org/>

⁸<https://onnx.ai/>

⁹<https://github.com/probprog/pyprob>

¹⁰We are planning to provide PPX bindings for these PPLs in future work.

then discarded, of the order $\sim 10^6$ for the Sherpa model we present in this paper. This is required to find the posterior density, which, as the model begins from an arbitrary point of the prior, can be very far from the starting region. Once this “burn-in” stage is completed the MCMC chain should be sampling from within the region containing the posterior. In addition to this, the sequential nature of each chain limits our ability to parallelize the computation, again creating computational inefficiencies in the high-dimensional space of simulator execution traces that we work with in our technique.

In order to provide fast, repeated inference in a distributed setting, we implement the IC algorithm, which trains a deep recurrent NN to provide proposals for an IS scheme [47]. This works by running the simulator many times and therefore sampling a large set of execution traces from the simulator prior $p(\mathbf{x}, \mathbf{y})$, and using these to train a NN that represents $q(\mathbf{x}|\mathbf{y})$, i.e., informed proposals for random number draws $\mathbf{x}$ given observations $\mathbf{y}$, by optimizing the loss $\mathcal{L}(\phi) = \mathbb{E}_{p(\mathbf{y})} [D_{\text{KL}}(p(\mathbf{x}|\mathbf{y})||q_\phi(\mathbf{x}|\mathbf{y}))] = \mathbb{E}_{p(\mathbf{x}, \mathbf{y})} [-\log q_\phi(\mathbf{x}|\mathbf{y})] + \text{const.}$, where ϕ are NN parameters (Algorithm 1 and Figure 3) [47]. This phase of sampling the training data and training the NN is costly, but it needs to be performed only once for any given model. Once the proposal NN is trained to convergence, the IC inference engine becomes competitive in performance, which allows us to achieve a given effective sample size in the posterior $p(\mathbf{x}|\mathbf{y})$ using a fraction of the RMH computational cost. IC inference is embarrassingly parallel, where many instances of the same trained NN can be executed to run distributed inference on a given observation.

To further improve inference performance, we make several low-level improvements in the code base. The C++ front end of PPX uses concatenated stack frames of each random number draw as a unique address identifying a latent variable in the corresponding PPL model. Stack traces are obtained with the `backtrace(3)` function as instruction addresses and then converted to symbolic names using the `dldaddr(3)` function [50]. The conversion is quite expensive, which prompted us to add a hash map to cache `dldaddr` results, giving a 5x improvement in the production of address strings that are essential in our inference engines. The particle detector simulator that we use was initially coded to use the `xtensor` library¹¹ to implement the probability density function (PDF) of multivariate normal distributions in the general case, but was exclusively called on 3D data. This code was replaced by a scalar-based implementation limited to the 3D case, resulting in a 13x speed-up in the PDF, and a 1.5x speed-up of our simulator pipeline in general. The bulk of our further optimizations focus on the NN training for IC inference and are discussed in the next sections.

4.3 Dynamic neural network architecture

The NN architecture used in IC inference is based on a LSTM [35] recurrent core that gets executed as many time steps as the simulator’s probabilistic trace length (Figure 3). To this core NN, various other NN components get attached according to the series of addresses A_t executed in the simulator. In other words, we construct a dynamic NN whose runtime structure changes in each execution trace, implemented using the dynamic computation graph infrastructure in PyTorch. The input to this LSTM in each time step is a concatenation of embeddings of the observation, the current address in

Figure 2: Loss curves for NN architectures considered in the hyperparameter search detailed in the text.

the simulator, and the previously sampled value. The observation embedding is a NN specific to the observation domain. Address embeddings are learned vectors representing the identity of random choices A_t in the simulator address space. Sample embeddings are address-specific layers encoding the value of the random draw in the previous time step. The LSTM output, at each time step, is fed into address-specific proposal layers that provide the final output of the NN for IC inference: proposal distributions $q(\mathbf{x}|\mathbf{y})$ to use for each address A_t as the simulator keeps running and requesting new random numbers over the PPX protocol (Section 4.1).

For the Sherpa experiments reported in this paper, we work with 3D observations of size $35 \times 35 \times 20$, representing particle detector voxels. To tune NN architecture hyperparameters, we search a grid of LSTM stacks in range $\{1, 4\}$, LSTM hidden units in the set $\{128, 256, 512\}$, and number of proposal mixture components in the set $\{5, 10, 25, 50\}$ (Figure 2). We settle on the following architecture: an LSTM with 512 hidden units; an observation embedding of size 256, encoded with a 3D convolutional neural network (CNN) [48] acting as a feature extractor, with layer configuration `Conv3D(1, 64, 3)–Conv3D(64, 64, 3)–MaxPool3D(2)–Conv3D(64, 128, 3)–Conv3D(128, 128, 3)–Conv3D(128, 128, 3)–MaxPool3D(2)–FC(2048, 256)`; previous sample embeddings of size 4 given by single-layer NNs; and address embeddings of size 64. The proposal layers are two-layer NNs, the output of which are either a mixture of ten truncated normal distributions [12] (for uniform continuous priors) or a categorical distribution (for categorical priors). We use ReLU nonlinearities in all NN components. All of these NN components except the LSTM and the 3DCNN are dependent on addresses A_t in the simulator, and these address-specific layers are created at the first encounter with a random number draw at a given address. Thus the number of trainable parameters in an IC NN is dependent on the size of the training data, because the more data gets used, the more likely it becomes to encounter new addresses in the simulator.

The pyprob framework is capable of operating in an “online” fashion, where NN training and layer generation happens using traces sampled by executing the simulator on-the-fly and discarding traces after each minibatch, or “offline”, where traces are sampled from the simulator and saved to disk as a dataset for further reuse (Algorithm 2). In our experiments, we used training datasets of 3M and 15M traces, resulting in NN sizes of 156,960,440 and 171,732,688

¹¹<https://xtensor.readthedocs.io>

Figure 3: Simulation and inference. Top: model addresses, priors and samples. Bottom: IC inference engine proposals and NN architecture.

parameters respectively. All timing and scaling results presented in Sections 6.1 and 6.2 are performed with the larger network.

4.4 Training of dynamic neural networks

Scalable training of dynamic NNs we introduced in Section 4.3 pose unique challenges. Because of the address-dependent nature of the embedding and proposal layers of the overall IC NN, different nodes/ranks in a distributed training setting will work with different NN configurations according to the minibatch of training data they process at any given time. When the same NN is not shared across all nodes/ranks, it is not possible to rely on a generic allreduce operation for gradient averaging which is required for multi-node synchronous SGD. Inspired by neural machine translation (NMT) [75], in the offline training mode with training data saved on the disk, we implemented the option of pre-processing the whole dataset to pre-generate all embedding and proposal layers that a given dataset would imply to exist. Once layer pre-generation is done, the collection of all embedding and proposal layers are shared on each node/rank. In this way, for offline training, we have a globally shared NN representing the superset of all NN components each node needs to handle in any given minibatch, thus making it possible to scale training of the NN on multiple nodes.

Our allreduce-based training algorithm can also work in the online training setting, where training data is sampled from the simulator on-the-fly, if we freeze a globally shared NN and discard any subsequently encountered traces that contain addresses unknown at the time of NN architecture freezing. In future work, we intend to add a distributed open-ended implementation for online training to allow running without discarding, that will require the NN instances in each node/rank to grow with newly seen addresses.

Algorithm 1 Computing minibatch loss $\mathcal{L}_n$ of NN parameters ϕ

Require: Minibatch $\mathcal{D}_n$
 $L \leftarrow$ number of unique trace types found in $\mathcal{D}_n$
Construct sub-minibatches $\mathcal{D}_n^l$, for $l = 1, \dots, L$
 $\mathcal{L}_n \leftarrow 0$
for $l \in \{1, \dots, L\}$ **do**
 $\mathcal{L}_n \leftarrow \mathcal{L}_n - \sum_{(x, y) \in \mathcal{D}_n^l} \log q_\phi(x|y)$
end for
return $\mathcal{L}_n$

Algorithm 2 Distributed training with MPI backend. $p(x, y)$ is the simulator and $\hat{G}(x, y)$ is an offline dataset sampled from $p(x, y)$

Require: OnlineData {True/False value}
Require: B {Minibatch size}
Initialize inference network $q_\phi(x|y)$
 $N \leftarrow$ number of processes
for all $n \in \{1, \dots, N\}$ **do**
while Not Stop **do**
if OnlineData **then**
Sample $\mathcal{D}_n = \{(x, y)_1, \dots, (x, y)_B\}$ from $p(x, y)$
else
Get $\mathcal{D}_n = \{(x, y)_1, \dots, (x, y)_B\}$ from $\hat{G}(x, y)$
end if
Synchronize parameters (ϕ) across all processes
 $\mathcal{L}_n \leftarrow -\frac{1}{B} \sum_{(x, y) \in \mathcal{D}_n} \log q_\phi(x|y)$
Calculate $\nabla_\phi \mathcal{L}_n$
Call all_reduce s.t. $\nabla_\phi \mathcal{L} \leftarrow \frac{1}{N} \sum_{n=1}^N \nabla_\phi \mathcal{L}_n$
Update ϕ using $\nabla_\phi \mathcal{L}$ with e.g. ADAM, SGD, LARC, etc.
end while
end for

4.4.1 Single node improvements to Etalumis. We profiled the Etalumis architecture with vtune, Cprofiler, and PyTorch autograd profiler, identifying data loading and 3D convolution as the primary computational hot-spots on which we focused our optimization efforts. We provide details on data loading and 3D convolution in the subsequent sections. In addition to these, execution traces from the Sherpa simulator have many different trace types (a unique sequence of addresses A_t , with different sampled values) with different rates of occurrence: in a given dataset, some trace types can be encountered thousands of times while others are seen only once. This is problematic because at training time we further divide each minibatch into “sub-minibatches” based on trace type, where each sub-minibatch can be processed by the NN in a single forward execution due to all traces being of the same type, i.e., sharing the same sequence of addresses A_t and therefore requiring the same NN structure (Algorithm 1). Therefore minibatches containing more than one trace type do not allow for effective parallelization and vectorization. In other words, unlike conventional NNs, the effective minibatch size is determined by the average size of sub-minibatches, and the more trace types we have within a minibatch, the slower the computation. To address this, we explored multiple methods to enlarge effective minibatch size, such as sorting traces, multi-bucketing, and selectively batching traces from the same trace type together in each minibatch. These options and their trade offs are described in more detail in Section 7.

4.4.2 Single node improvements to PyTorch. The flexibility of dynamic computation graphs and competitive speed of PyTorch have been crucial for this project. Optimizations were performed on code belonging to PyTorch stable release v1.0.0 to better support this project on Intel® Xeon® CPU platforms, focused in particular on 3D convolution operations making use of the MKL-DNN open source math library. MKL-DNN uses a direct convolution algorithm and for a 5-dimensional input tensor with layout {N, C, D, H, W}, it is reordered into a layout of {N, C, D, H, W, 8c} which is more amenable for SIMD vectorization.¹² The 3D convolution operator is vectorized on the innermost dimension which matches the 256-bit instruction length on AVX2, and parallelized on the outer dimensions. We also performed cache optimization to further improve performance. With these improvements we found the heavily used 3D convolution kernel achieved an 8x improvement on the Cori HSW platform.¹³ The overall improvement on single node training time is given in Section 6.1. These improvements are made available in a fork of the official PyTorch repository.¹⁴

4.4.3 I/O optimization. I/O is challenging in many deep learning workloads partly due to random access patterns, such as those induced by shuffling, disturbing any pre-determined access order. In order to reduce the number of random access I/O operations, we developed a parallel trace sorting algorithm and pre-sorted the 15M traces according to trace type (Section 4.4.1). We further grouped the small trace files into larger files, going from 750 files with 20k traces per file to 150 files with 100k traces per file. With this grouping and sorting, we ensured that I/O requests follow a sequential access onto a contiguous file region which further improved the I/O performance. Metadata operations are also costly, so we enhanced the Python shelve module’s file open/close performance with a caching mechanism, which allows concurrent access from different ranks to the same file.

Specific to our PPL setting, training data consists of execution traces that have a complex hierarchy, with each trace file containing many trace objects that consist of variable sequences of sample objects representing random number draws, which further contain variable length tensors, strings, integers, booleans, and other generic Python objects. PyTorch serialization with pickle is used to handle the complex trace data structure, but the pickle and unpickle overhead are very high. We developed a “pruning” function to shrink the data by removing non-necessary structures. We also designed a dictionary of simulator addresses A_t , which accumulates the fairly long address strings and assigns shorthand IDs that are used in serialization. This brought a 40% memory consumption reduction as well as large disk space saving.

For distributed training, we developed distributed minibatch sampler and dataset classes conforming to the PyTorch training API. The sampler first splits the sorted trace indices into minibatch-sized chunks, so that all traces in each minibatch are highly likely to be of the same type, then optionally groups these chunks into several buckets (Section 7.2). Within each bucket, the chunks are assigned with a round-robin algorithm to different ranks, such

that each rank has roughly same distribution of workload. The distributed sampler enables us to scale the training at 1,024 nodes.

The sorting of traces and their grouping into minibatch chunks significantly improves the training speed (up to 50x in our experiments) by enabling all traces in a minibatch to be propagated through the NN in the same forward execution, in other words, decreasing the need for “sub-minibatching” (Section 4.4.1). This sorting and chunking scheme generates minibatches that predominantly contain a single trace type. However, the minibatches used at each iteration are sampled randomly without replacement from different regions of the sorted training set, and therefore contain different trace types, resulting in a gradient unbiased in expectation during any given epoch.

In our initial profiling, the cost of I/O was more than 50% of total run time. With these data re-structuring and parallel I/O optimizations, we reduced the I/O to less than 5%, achieving 10x speedup at different scales.

4.4.4 Distributed improvements to PyTorch MPI CPU code. PyTorch has a torch.distributed backend,¹⁵ which allows scalable distributed training with high performance on both CPU and GPU clusters. Etalumis uses the MPI backend as appropriate for the synchronous SGD setting that we implement (Algorithm 2) and the HPC machines we utilize (Section 5). We have made various improvements to this backend to enable the large-scale distributed training on CPU systems required for this project. The call torch.distributed.all_reduce is used to combine the gradient tensors for all distributed MPI ranks. In Etalumis, the set of non-null gradient tensors differs for each rank and is a small fraction of the total set of tensors. Therefore we first perform an allreduce to obtain a map of all the tensors that are present on all ranks; then we create a list of the tensors, filling in the ones that are not present on our rank with zero; finally, we reduce all of the gradient tensors in the list. PyTorch all_reduce does not take a list of tensors so normally a list comprehension is used, but this results in one call to MPI_Allreduce for each tensor. We modified PyTorch all_reduce to accept a list of tensors. Then, in the PyTorch C++ code for allreduce, we concatenate small tensors into a buffer, call MPI_Allreduce on the buffer, and copy the results back to the original tensor. This eliminates almost all the allreduce latency and makes the communication bandwidth-bound.

We found that changing Etalumis to reduce only the non-null gradients gives a 4x improvement in allreduce time. Tensor concatenation improves overall performance by an additional 4% on one node which increases as nodes are added. With these improvements, the load balance effects discussed in Sections 6.2 and 7.2 are dominant and so are our primary focus of further distributed optimizations. Other future work could include performing the above steps for each backward layer with an asynchronous allreduce to overlap the communications for the previous layer with the computation for the current layer.

5 SYSTEMS AND SOFTWARE

5.1 Cori

We use the “data” partition of the Cori system at the National Energy Research Scientific Computing Center (NERSC) at Lawrence

¹²https://intel.github.io/mkl-dnn/understanding_memory_formats.html

¹³<https://docs.nersc.gov/analytics/machinelearning/benchmarks/>

¹⁴Intel-optimized PyTorch: <https://github.com/intel/pytorch>

¹⁵<https://pytorch.org/docs/stable/distributed.html>

Table 1: Intel® Xeon® CPU models and codes

Model	Code
E5-2695 v2 @ 2.40GHz (12 cores/socket)	IVB
E5-2698 v3 @ 2.30GHz (16 cores/socket)	HSW
E5-2697A v4 @ 2.60GHz (16 cores/socket)	BDW
Platinum 8170 @ 2.10GHz (26 cores/socket)	SKL
Gold 6252 @ 2.10GHz (24 cores/socket)	CSL

Berkeley National Laboratory. Cori is a Cray XC40 system, and the data partition features 2,388 nodes. Each node has two sockets and each socket is populated with a 16-core 2.3 GHz Intel® Xeon® E5-2698 v3 CPU (referred to as HSW from now on), with peak single-precision (SP) performance of 1.2 Tflop/s and 128 GB of DDR4-2133 DRAM. Nodes are connected via the Cray Aries low-latency, high-bandwidth interconnect utilizing the dragonfly topology. In addition, the Cori system contains 288 Cray DataWarp nodes (also known as the “Burst Buffer”) which house the input datasets for the Cori experiments presented here. Each DataWarp node contains 2×3.2 TB SSDs, giving a system total of 1.8 PB of SSD storage, with up to 1.7 TB/sec read/write performance and over 28M IOP/s. Cori also has a Sonnet 2000 Lustre filesystem, which consists of 248 Object Storage Targets (OSTs) and 10,168 disks, giving nearly 30 PB of storage and a maximum of 700 GB/sec IO performance. This filesystem is used for output files (networks and logs) for Cori experiments and both input and output for the Edison experiments.

5.2 Edison

We also make use of the Edison system at NERSC. Edison is a Cray XC30 system with 5,586 nodes. Each node has two sockets, each socket is populated with a 12-core 2.4 GHz Intel® Xeon® E5-2695 v2 CPU (referred to as IVB from now on), with peak performance of 460.8 SP Gflop/s, and 64 GB DDR3-1866 memory. Edison mounts the Cori Lustre filesystem described above.

5.3 Diamond cluster

In order to evaluate and improve the performance on newer Intel® processors we make use of the Diamond cluster, a small heterogeneous cluster maintained by Intel Corporation. The interconnect uses Intel® Omni-Path Architecture switches and host adapters. The nodes used for the results in this paper are all two socket nodes. Table 1 presents the CPU models used and the three letter abbreviations used in this paper.

5.4 Particle physics simulation software

In our experiments we use Sherpa version 2.2.3, coupled to a fast 3D detector simulator that we configure to use $20 \times 35 \times 35$ voxels. Sherpa is implemented in C++, and therefore we use the C++ front end for PPX. We couple to Sherpa by a system-wide rerouting of the calls to the random number generator, which is made easy by the existence of a third-party random number generator interface (External_RNG) already present in Sherpa.

For this paper, in order to facilitate reproducible experiments, we run in the offline training mode and produce a sample of 15M traces that occupy 1.7 TB on disk. Generation of this 15M dataset

Table 2: Single node training throughput in traces/sec and flop rate (Gflop/s). 1-socket throughput and flop rate are for a single process while 2-socket is for 2 MPI processes on a single node.

Platform	1-socket traces/s	2-socket traces/s	1-socket Gflop/s (% peak)
IVB (Edison)	13.9	25.6	196 (43%)
HSW (Cori)	32.1	56.5	453 (38%)
BDW (Diamond)	30.5	57.8	430 (32%)
SKL (Diamond)	49.9	82.7	704 (20%)
CSL (Diamond)	51.1	93.1	720 (22%)

was completed in 3 hours on 32 IVB nodes of Edison. The traces are stored using Python `shelve`¹⁶ serialization, allowing random access to all entries contained in the collection of files with 100k traces in each file. These serialized files are accessed via the Python `dbm` module using the `gdbm` backend.

6 EXPERIMENTS AND RESULTS

6.1 Single node performance

We ran single node tests with one rank per socket for one and two ranks on the IVB nodes on Edison, the HSW partition of Cori and the BDW, SKL, and CSL nodes of the Diamond cluster. Table 2 shows the throughput and single socket flop rate and percentage of peak theoretical flop rate. We find that the optimizations described in Section 4.4.2 provide an improvement of 7x on the overall single socket run throughput (measured on HSW) relative to a default PyTorch version v1.0.0 installed via the official conda channel. We achieve **430 SP Gflop/s** on a single socket of the BDW system, measured using the available hardware counters for 256-bit packed SIMD single precision operations. This includes IO and is averaged over an entire 300k trace run. This can be compared to a theoretical peak flop rate for that BDW socket of 1,331 SP Gflop/s. Flop rates for other platforms are scaled from this measurement and given in Table 2. For further profiling we instrument the code with timers for each phase of the training (in order): minibatch read, forward, backward, and optimize. Figure 4 shows a breakdown of the time spent on a single socket after the optimizations described in Sections 4.4.1 and 4.4.2.¹⁷

6.2 Multi-node performance

In addition to the single socket operations, we time the two synchronization (allreduce) phases (gradient and loss). This information is recorded for each rank and each minibatch. Postprocessing finds the rank with the maximum work time (sum of the four phases mentioned in Section 6.1) and adds the times together. This gives the actual execution time. Further, we compute the *average* time across ranks for each work phase for each minibatch and add those together, giving the best time assuming no load imbalance. The results are shown in Figure 4. Comparing single socket results with 2 and 64 socket results shows the increased impact of load imbalance

¹⁶<https://docs.python.org/3/library/shelve.html>

¹⁷See disclaimers section after conclusions.

Figure 4: Actual and estimated best times for 1, 2, and 64 sockets. Horizontal bars at the top are to aid comparison between columns.

Figure 5: Mean loss and standard deviation (shaded) for five experiments with 128k minibatch size.

as nodes are added. This demonstrates a particular challenge of this project where the load on each node depends on the random minibatch of traces sampled for training on that node. The challenge and possible mitigation approaches we explored are discussed further in Section 7.2.

6.3 Large scale training on Cori and Edison

In order to choose training hyperparameters, we explored global minibatch sizes of {64, 2k, 32k, 128k}, and learning rates in the range $[10^{-7}, 10^{-1}]$ with grid search and compared the loss values after one epoch to find the optimal learning rate for different global minibatch sizes separately. For global minibatch sizes of 2k, 32k, and 128k, we trained with both Adam and Adam-LARC optimizers and compared loss value changes as a function of iterations. For training

at 1,024 nodes we choose to use 32k and 128k global minibatch sizes. For the 128k minibatch size, best convergence was found using the Adam-LARC optimizer with a polynomial decay (order=2) learning rate schedule [76] that decays from an initial global learning rate of 5.70×10^{-4} to final 2×10^{-5} after completing 12 epochs for the dataset with 15M traces. In Figure 5, we show the mean and standard-deviation for five training runs with this 128k minibatch size and optimizer, demonstrating stable convergence.

Figure 6 shows weak scaling results obtained for distributed training to over a thousand nodes on both the Cori and Edison systems. We use a fixed local minibatch size of 64 per rank with 2 ranks per node, and plot the mean and standard deviation throughput for each iteration in terms of traces/s (labeled “average” in the plot). We also show the fastest iteration (labeled “peak”). The average scaling efficiency at 1,024 nodes is 0.79 on Edison and 0.5 on Cori. The throughput at 1,024 nodes on Cori and Edison is 28,000 and 22,000 traces/s on average, with the peak as 42,000 and 28,000 traces/s respectively. One can also see that there is some variation in this performance due to the different compute times taken to process execution traces of different length and the related load imbalance as discussed in Sections 6.2 and 7.2. We determine the maximum sustained performance over a 10-iteration sliding window to be **450 Tflop/s** on Cori and **325 Tflop/s** on Edison.¹⁸

We have performed distributed training with global minibatch sizes of 32k and 128k at 1,024-node scale for extended periods to achieve convergence on both Cori and Edison systems. This is illustrated in Figure 7 where we show the loss for training and validation datasets as a function of iteration for an example run on Edison.

6.4 Inference and science results

Using our framework and the NNs trained using distributed resources at NERSC as described previously, we perform inference on test τ observation data that has not been used for training. As the approach of applying probabilistic programming in the setting of large-scale existing simulators is completely novel, there is no direct baseline in literature that provides the full posterior in each of these variables. Therefore we use our own MCMC (RMH)-based posterior as a baseline for the validation of the IC approach. We establish the convergence of the RMH posterior by running two independent MCMC chains with different initializations and computing the the Gelman–Rubin convergence metric [24] to confirm that they converge onto the same posterior distribution (Section 4.2).

Figure 8 shows a comparison of inference results from the RMH and IC approaches. We show selected latent variables (addresses) that are a small subset of the more than 24k addresses that were encountered in the prior space of the Sherpa experimental setup, but are of physics interest in that they correspond to properties of the τ particle. It can be seen that there is close agreement between the RMH and IC posterior distributions validating that our network has been adequately trained. We have made various improvements to the RMH inference processing rate but this form of inference is compute intensive and takes **115 hours** on a Edison IVB node to produce the 7.68M trace result shown. The corresponding 2M trace IC result completed in **30 mins** (achieving a 230× speedup

¹⁸See disclaimers section after conclusions.

Figure 6: Weak scaling on Cori and Edison systems, showing throughput for different node counts with a fixed local minibatch size of 64 traces per MPI rank with 2 ranks per node. Average (mean) over all iterations and the peak single iteration are shown. Ideal scaling is derived from the mean single-node rate with a shaded uncertainty from the standard deviation in that rate.

Figure 7: Training and validation loss for a 128k minibatch size experiment with the configuration described in the text run on 1,024 nodes of the Edison system.

for a comparable posterior result) on 24 HSW nodes, enabled by the parallelism of IC inference.

In addition to parallelization, a significant advantage of the IC approach is that it is amortized. This means that once the proposal NN is trained for any given model, it can be readily applied to large volumes of new collision data. Moreover IC inference runs with high effective sample sizes in comparison to RMH: each sample from the IC NN is an independent sample from the proposal distribution, which approaches the true posterior distribution with more training, whereas our autocorrelation measurements in the RMH posterior indicate that a very large number of iterations are needed to get statistically independent traces (on the order of $\sim 10^5$ for the type of decay event we use as the observation). These features of IC inference combine to provide a tractable approach for fast

Bayesian inference in complex models implemented by large-scale simulators.

7 DISCUSSION

The dynamic NN architecture employed for this project has presented a number of unique challenges for distributed training, which we covered in Section 4.4. Our innovations proved successful in enabling the training of this architecture at scale, and in this section we capture some of the lessons learned and unresolved issues encountered in our experiments.

7.1 Time to solution: trade-off between throughput and convergence

7.1.1 Increasing effective local minibatch size. As mentioned in Section 4.4.1, the distributed SGD scheme given in Algorithms 1 and 2 uses random traces sampled from the simulator, and can suffer from slow training throughput if computation cannot be efficiently parallelized across the full minibatch due to the presence of different trace types. Therefore, we explored the following methods to improve effective minibatch size: sorting the traces before batching, online batching of the same trace types together, and multi-bucketing. Each of these methods can improve throughput, but risk introducing bias into SGD training or increasing the number of iterations to converge, so we compare the wall clock time for convergence to determine the relative trade-off. The impact of multi-bucketing is described and discussed in section 7.2 below. Sorting and batching traces from the same trace type together offers considerable throughput increase with a relatively small impact on convergence, when combined with shuffling of minibatches for randomness, so we used these throughput optimizations for the results in this paper.

7.1.2 Choice of optimizers, learning rate scaling and scheduling for convergence. There is considerable recent literature on techniques

Figure 8: A comparison of posterior distributions obtained by RMH (filled histograms) and IC (outline histograms) and the ground truth values (dashed vertical lines) for a test τ decay observation. We show an illustrative subset of the latent variables, including x , y and z components of the τ -lepton momentum (top row), the energies of the two highest energy particles produced by the τ decay (middle left and bottom center), a contour plot showing correlation between these (bottom left), the τ decay channel ($\tau \rightarrow \pi \nu_\tau$ as mode) (middle right), and the missing transverse energy (bottom right).

for distributed and large-minibatch training, including scaling learning rates with number of nodes [31], alternative optimizers for large-minibatch-size training, and learning rate decay [27, 41, 76]. The network employed in this project presents a very different use-case to those considered in literature, prompting us to document our experiences here. For learning rate scaling with node count, we found sub-sqrt learning rate scaling works better than linear for an Adam-based optimizer [41]. We also compared Adam with the newer Adam-LARC optimizer [27, 76] for convergence performance and found the Adam-LARC optimizer to perform better for the very large global minibatch size 128k in our case. For smaller global minibatch sizes of 32k or lower, both plain Adam and Adam-LARC performed equally well. Finally, for the learning rate decay scheduler, we explored the following options: no decay, multi-step decay (per epoch), and polynomial decay of order 1 or 2 (calculated per iteration) [76]. We found that learning-rate decay can improve training performance and polynomial decay of order 2 provided the most effective schedule.

7.2 Load balancing

As indicated in Section 6.2, our work involves distinct scaling challenges due to the variation in compute time depending on execution trace length, address-dependent proposal and embedding layers, and representation of trace types inside each minibatch. These factors contribute to load imbalance. The trace length variation bears similarity to varying sequence lengths in NMT; however, unlike that case it is not possible to truncate the execution traces, and padding would introduce a cost to overall number of operations.

To resolve this load imbalance issue, we have explored a number of options, building on those from NMT, including a *multi-bucketing* scheme and a novel *dynamic batching* approach.

In *multi-bucketing* [10, 20, 40], traces are grouped into several buckets based on lengths, and every global minibatch is solely taken from a randomly picked bucket for every iteration. Multi-bucketing not only helps to balance the load among ranks for the same iteration, but also increases the effective minibatch size as traces from the same trace type have a higher chance to be in the same minibatch than in the non-bucketing case, achieving higher throughput. For a local minibatch-size of 16 with 10 buckets we measured throughput increases in the range of 30–60% at 128–256 node scale on Cori. However, our current multi-bucketing implementation does multiple updates in the same bucket continuously. When this implementation is used together with batching the traces from the same trace type together (as discussed above in Section 7.1.1), it negatively impacts convergence behavior. We believe this to be due to the fact that training on each specific bucket for multiple updates introduces over-fitting onto that specific subset of networks so moving to a new bucket for multiple updates causes information of the progress made with previous buckets to be lost. As such we did not employ this configuration for the results reported in this paper.

With *dynamic batching*, we replaced the requirement of fixed minibatch size per rank with a desired number of “tokens” per rank (where a token is a unit of random number draws in each trace), so that we can, for instance, allocate many short traces (with smaller number of tokens each) for one rank but only a few long traces for another rank, in order to balance the load for the LSTM network due to length variation. While an equal-token approach has been used in NMT, this did not offer throughput gains for our model, which has an additional 3DCNN component in which the compute time depends on the number of traces within the local minibatch, so if dynamic batching only considers total tokens per rank for the LSTM it can negatively impact the 3DCNN load.

Through these experiments we found that our current optimal throughput and convergence performance came from not employing these load-balancing schemes although we intend to explore modifications to these schemes as ongoing work.

8 SCIENCE IMPLICATIONS AND OUTLOOK

We have provided a common interface to connect PPLs with simulators written in arbitrary code in a broad range of programming languages. This opens up possibilities for future work in all applied fields where simulators are used to model real-world systems, including epidemiology modeling such as disease transmission and prevention models [66], autonomous vehicle and reinforcement learning environments [21], cosmology [7], and climate science

[68]. In particular, the ability to control existing simulators at scale and to generate interpretable posteriors is relevant to scientific domains where interpretability in model inference is critical.

We have demonstrated both MCMC- and IC-based inference of detector data originating from τ -decays simulated with the Sherpa Monte Carlo generator at scale. This offers, for the first time, the potential of Bayesian inference on the full latent structure of the large numbers of collision events produced at accelerators such as the LHC, enabling deep interpretation of observed events. For instance, ambiguity in the decay of a particle can be related exactly to the physics processes in the simulator that would give rise to that ambiguity. In order to fully realize this potential, future work will expand this framework to more complex particle decays (such as the Higgs decay to τ leptons) and incorporate a more detailed detector simulation (e.g., Geant4 [6]). We will demonstrate this on a full LHC physics analysis, reproducing the efficiency of point-estimates, together with the full posterior and interpretability, so that this can be exploited for discovery of new fundamental physics.

The IC objective is designed so that the NN proposal $q(\mathbf{x}|\mathbf{y})$ approximates the posterior $p(\mathbf{x}|\mathbf{y})$ asymptotically closely with more training. This costly training phase needs to be done *only once* for a given simulator-based model, giving us a NN that can provide samples from the model posterior in parallel for any new observed data. In this setting where we have a fast, amortized $q(\mathbf{x}|\mathbf{y}) \approx p(\mathbf{x}|\mathbf{y})$, our ultimate goal is to add the machinery of Bayesian inference to the toolbox for critical tasks such as triggering [1] and event reconstruction by conditioning on potentially interesting events (e.g., $q(\text{ParticleType}|\cdot) \geq \epsilon$). Recent activity exploring the use of FPGAs for NN inference for particle physics [23] will help implementation of these approaches, and HPC systems will be crucial in the training and inference phases of such frameworks.

9 CONCLUSIONS

Inference in simulator-based models remains a challenging problem with potential impact across many disciplines [14, 60]. In this paper we present the first probabilistic programming implementation capable of controlling existing simulators and running at large-scale on HPC platforms. Through the PPX protocol, our framework successfully couples with large-scale scientific simulators leveraging thousands of lines of existing simulation code encoding domain-expert knowledge. To perform efficient inference we make use of the inference compilation technique, and we train a dynamic neural network involving LSTM and 3DCNN components, with a large global minibatch size of 128k. IC inference achieved a 230 $\times$ speedup compared with the MCMC baseline. We optimize the popular PyTorch framework to achieve a significant single-socket speedup for our network and 20–43% of peak theoretical flop rate on a range of current CPUs.¹⁹ We augment and develop PyTorch’s MPI implementation to run it at the unprecedented scale of 1,024 nodes (32,768 and 24,576 cores) of the Cori and Edison supercomputers with a sustained flop rate of 0.45 Pflop/s. We demonstrate we can successfully train this network to convergence at these large scales, and use this to perform efficient inference on LHC collision events. The developments described here open the door for exploiting HPC

resources and existing detailed scientific simulators to perform rapid Bayesian inference in very complex scientific settings.

DISCLAIMERS

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

For more complete information visit www.intel.com/benchmarks.

Performance results are based on testing as of March 22 and March 28, 2019 and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure.

Configurations: Testing on Cori and Edison (see §6.1) was performed by NERSC. Testing on Diamond cluster was performed by Intel.

Intel technologies’ features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. Check with your system manufacturer or retailer or learn more at www.intel.com.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Intel, VTune, Xeon are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

ACKNOWLEDGMENTS

The authors would like to acknowledge valuable discussions with Thorsten Kurth on scaling aspects, Quincey Koziol on I/O; Steve Farrell on NERSC PyTorch, Holger Schulz on Sherpa, and Xiaoming Cui, from Intel AIPG team, on NMT. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility operated under Contract No. DE-AC02-05CH11231. This work was partially supported by the NERSC Big Data Center; we acknowledge Intel for their funding support. KC, LH, and GL were supported by the National Science Foundation under the awards ACI-1450310. Additionally, KC was supported by the National Science Foundation award OAC-1836650. BGH is supported by the EPSRC Autonomous Intelligent Machines and Systems grant. AGB and PT are supported by EPSRC/MURI grant EP/N019474/1 and AGB is also supported by Lawrence Berkeley National Lab. FW is supported by DARPA D3M, under Cooperative Agreement FA8750-17-2-0093, Intel under its LBNL NERSC Big Data Center, and an NSERC Discovery grant.

REFERENCES

- [1] Morad Aaboud et al. 2017. Performance of the ATLAS Trigger System in 2015. *European Physical Journal C* 77, 5 (2017), 317. <https://doi.org/10.1140/epjc/s10052-017-4852-3>
- [2] G. Aad et al. 2008. The ATLAS Experiment at the CERN Large Hadron Collider. *JINST* 3 (2008), S08003. <https://doi.org/10.1088/1748-0221/3/08/S08003>
- [3] G. Aad et al. 2015. Search for the Standard Model Higgs boson produced in association with top quarks and decaying into $b\bar{b}$ in pp collisions at $\sqrt{s}=8$ TeV with the ATLAS detector. *The European Physical Journal C* 75, 7 (29 Jul 2015), 349. <https://doi.org/10.1140/epjc/s10052-015-3543-1>
- [4] G. Aad et al. 2016. Reconstruction of hadronic decay products of tau leptons with the ATLAS experiment. *The European Physical Journal C* 76, 5 (25 May 2016), 295. <https://doi.org/10.1140/epjc/s10052-016-4110-0>
- [5] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. Tensorflow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*. 265–283.
- [6] Sea Agostinelli, John Allison, K al Amako, J Apostolakis, H Araujo, P Arce, M Asai, D Axen, S Banerjee, G Barrand, et al. 2003. GEANT4—a simulation toolkit. *Nuclear instruments and methods in physics research section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506, 3 (2003), 250–303.
- [7] Joël Akeret, Alexandre Refregier, Adam Amara, Sebastian Seehars, and Caspar Hasner. 2015. Approximate Bayesian computation for forward modeling in cosmology. *Journal of Cosmology and Astroparticle Physics* 2015, 08 (2015), 043.
- [8] M Sanjeev Arulampalam, Simon Maskell, Neil Gordon, and Tim Clapp. 2002. A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions on Signal Processing* 50, 2 (2002), 174–188.
- [9] Atılım Güneş Baydin, Barak A. Pearlmutter, Alexey Andreyevich Radul, and Jeffrey Mark Siskind. 2018. Automatic differentiation in machine learning: a survey. *Journal of Machine Learning Research (JMLR)* 18, 153 (2018), 1–43. <http://jmlr.org/papers/v18/17-468.html>

¹⁹See disclaimers section after conclusions.

- [10] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. In *Proceedings of the 26th Annual International Conference on Machine Learning*. ACM, 41–48.
- [11] Eli Bingham, Jonathan P Chen, Martin Jankowiak, Fritz Obermeyer, Neeraj Pradhan, Theofanis Karaletsos, Rohit Singh, Paul Szerlip, Paul Horsfall, and Noah D Goodman. 2018. Pyro: Deep universal probabilistic programming. *Journal of Machine Learning Research* (2018).
- [12] Christopher M Bishop. 1994. *Mixture density networks*. Technical Report NCRG/94/004. Neural Computing Research Group, Aston University.
- [13] Christopher M Bishop. 2006. *Pattern Recognition and Machine Learning*. Springer.
- [14] Johann Brehmer, Gilles Louppe, Juan Pavez, and Kyle Cranmer. 2018. Mining gold from implicit models to improve likelihood-free inference. *arXiv preprint arXiv:1805.12244* (2018).
- [15] Jianmin Chen, Xinghao Pan, Rajat Monga, Samy Bengio, and Rafal Jozefowicz. 2016. Revisiting distributed synchronous SGD. *arXiv preprint arXiv:1604.00981* (2016).
- [16] D. Das, S. Avancha, D. Mudigere, K. Vaidynathan, S. Sridharan, D. Kalamkar, B. Kaul, and P. Dubey. 2016. Distributed Deep Learning Using Synchronous Stochastic Gradient Descent. *ArXiv e-prints* (Feb. 2016). arXiv:cs.DC/1602.06709
- [17] Jeffrey Dean, Greg S. Corrado, Rajat Monga, Kai Chen, Matthieu Devin, Quoc V. Le, Mark Z. Mao, Marc'Aurelio Ranzato, Andrew Senior, Paul Tucker, Ke Yang, and Andrew Y. Ng. 2012. Large Scale Distributed Deep Networks. In *Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 1 (NIPS'12)*. Curran Associates Inc., USA, 1223–1231. <http://dl.acm.org/citation.cfm?id=2999134.2999271>
- [18] Adji Bousso Dieng, Dustin Tran, Rajesh Ranganath, John Paisley, and David Blei. 2017. Variational Inference via χ Upper Bound Minimization. In *Advances in Neural Information Processing Systems*. 2732–2741.
- [19] Joshua V Dillon, Ian Langmore, Dustin Tran, Eugene Brevdo, Srinivas Vasudevan, Dave Moore, Brian Patton, Alex Alemi, Matt Hoffman, and Rif A Saurous. 2017. TensorFlow distributions. *arXiv preprint arXiv:1711.10604* (2017).
- [20] Patrick Doetsch, Pavel Golik, and Hermann Ney. 2017. A comprehensive study of batch construction strategies for recurrent neural networks in mxnet. *arXiv preprint arXiv:1705.02414* (2017).
- [21] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. 2017. CARLA: An Open Urban Driving Simulator. In *Proceedings of the 1st Annual Conference on Robot Learning*. 1–16.
- [22] Arnaud Doucet and Adam M Johansen. 2009. A tutorial on particle filtering and smoothing: Fifteen years later. (2009).
- [23] Javier Duarte et al. 2018. Fast inference of deep neural networks in FPGAs for particle physics. *JINST* 13, 07 (2018), P07027. <https://doi.org/10.1088/1748-0221/13/07/P07027> arXiv:physics.ins-det/1804.06913
- [24] Andrew Gelman, Hal S Stern, John B Carlin, David B Dunson, Aki Vehtari, and Donald B Rubin. 2013. *Bayesian data analysis*. Chapman and Hall/CRC.
- [25] Samuel Gershman and Noah Goodman. 2014. Amortized inference in probabilistic reasoning. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, Vol. 36.
- [26] Zoubin Ghahramani. 2015. Probabilistic machine learning and artificial intelligence. *Nature* 521, 7553 (2015), 452.
- [27] B. Ginsburg, I. Gitman, and O. Kuchaiev. 2018. Layer-Wise Adaptive Rate Control for Training of Deep Networks. in *preparation* (2018).
- [28] Sheldon L Glashow. 1961. Partial-symmetries of weak interactions. *Nuclear Physics* 22, 4 (1961), 579–588.
- [29] Tanju Gleisberg, Stefan Hölche, F Krauss, M Schönherr, S Schumann, F Siegert, and J Winter. 2009. Event generation with SHERPA 1.1. *Journal of High Energy Physics* 2009, 02 (2009), 007.
- [30] Prem Gopalan, Wei Hao, David M Blei, and John D Storey. 2016. Scaling probabilistic models of genetic variation to millions of humans. *Nature Genetics* 48, 12 (2016), 1587.
- [31] Priya Goyal, Piotr Dollar, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. 2017. Accurate, Large Minibatch SGD: Training ImageNet in 1 Hour. *arXiv preprint arXiv:1706.02677v1* (2017).
- [32] David Griffiths. 2008. *Introduction to elementary particles*. John Wiley & Sons.
- [33] Ralf Herbrich, Tom Minka, and Thore Graepel. 2007. TrueSkill™: a Bayesian skill rating system. In *Advances in Neural Information Processing Systems*. 569–576.
- [34] Pieter Hintjens. 2013. *ZeroMQ: messaging for many applications*. O'Reilly Media, Inc.
- [35] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation* 9, 8 (1997), 1735–1780.
- [36] Matthew D Hoffman, David M Blei, Chong Wang, and John Paisley. 2013. Stochastic variational inference. *The Journal of Machine Learning Research* 14, 1 (2013), 1303–1347.
- [37] Matthew D Hoffman and Andrew Gelman. 2014. The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research* 15, 1 (2014), 1593–1623.
- [38] F. N. Iandola, K. Ashraf, M. W. Moskewicz, and K. Keutzer. 2015. FireCaffe: near-linear acceleration of deep neural network training on compute clusters. *ArXiv e-prints* (Oct. 2015). arXiv:cs.CV/1511.00175
- [39] Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. 2016. On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836* (2016).
- [40] Viacheslav Khomenko, Oleg Shyshkov, Olga Radyvonenko, and Kostiantyn Bokhan. 2016. Accelerating recurrent neural network training using sequence bucketing and multi-gpu data parallelization. In *2016 IEEE First International Conference on Data Stream Mining & Processing (DSMP)*. IEEE, 100–103.
- [41] D. P. Kingma and J. Ba. 2014. Adam: A Method for Stochastic Optimization. *ArXiv e-prints* (Dec. 2014). arXiv:cs.LG/1412.6980
- [42] Diederik P Kingma and Max Welling. 2013. Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114* (2013).
- [43] Kunitaka Kondo. 1988. Dynamical likelihood method for reconstruction of events with missing momentum. I. Method and toy models. *Journal of the Physical Society of Japan* 57, 12 (1988), 4126–4140.
- [44] Thorsten Kurth, Sean Treichler, Joshua Romero, Mayur Mudigonda, Nathan Luehr, Everett Phillips, Ankur Mahesh, Michael Matheson, Jack Deslippe, Massimiliano Fatica, Prabhat, and Michael Houston. 2018. Exascale Deep Learning for Climate Analytics. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC '18)*. IEEE Press, Piscataway, NJ, USA, Article 51, 12 pages. <http://dl.acm.org/citation.cfm?id=3291656.3291724>
- [45] Thorsten Kurth, Jian Zhang, Nadathur Satish, Evan Racah, Ioannis Mitliagkas, Md Mostofa Ali Patwary, Tareq Malas, Narayanan Sundaram, Wahid Bhimji, Mikhail Smorkalov, et al. 2017. Deep learning at 15PF: supervised and semi-supervised classification for scientific data. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis. ArXiv e-prints*, 7. arXiv:1708.05256
- [46] Tuan Anh Le. 2015. Inference for higher order probabilistic programs. *Masters thesis, University of Oxford* (2015).
- [47] Tuan Anh Le, Atılım Güneş Baydin, and Frank Wood. 2017. Inference Compilation and Universal Probabilistic Programming. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS) (Proceedings of Machine Learning Research)*, Vol. 54. PMLR, Fort Lauderdale, FL, USA, 1338–1348.
- [48] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (1998), 2278–2324.
- [49] Mario Lezcano Casado, Atılım Güneş Baydin, David Martinez Rubio, Tuan Anh Le, Frank Wood, Lukas Heinrich, Gilles Louppe, Kyle Cranmer, Wahid Bhimji, Karen Ng, and Prabhat. 2017. Improvements to Inference Compilation for Probabilistic Programming in Large-Scale Scientific Simulators. In *Neural Information Processing Systems (NIPS) 2017 workshop on Deep Learning for Physical Sciences (DLPS)*, Long Beach, CA, US, December 8, 2017.
- [50] Linux man-pages project. 2019. *Linux Programmer's Manual*. <http://man7.org/linux/man-pages/index.html>
- [51] Amrita Mathuriya, Deborah Bard, Peter Mendenygral, Lawrence Meadows, James Armemann, Lei Shao, Siyu He, Tuomas Karna, Daina Moise, Simon J. Pennycook, Kristyn Maschoff, Jason Sewall, Nalini Kumar, Shirley Ho, Mike Ringenburg, Prabhat, and Victor Lee. 2018. CosmoFlow: Using Deep Learning to Learn the Universe at Scale. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC '18)*. IEEE Press, Piscataway, NJ, USA, Article 65, 11 pages. <https://dl.acm.org/citation.cfm?id=3291743>
- [52] Amrita Mathuriya, Thorsten Kurth, Vivek Rane, Mustafa Mustafa, Lei Shao, Debbie Bard, Victor W Lee, et al. 2017. Scaling GRPC Tensorflow on 512 nodes of Cori Supercomputer. In *Neural Information Processing Systems (NIPS) 2017 workshop on Deep Learning At Supercomputer Scale*, Long Beach, CA, US, December 8, 2017.
- [53] Sam McCandlish, Jared Kaplan, and et.al Amodei, Dario. 2018. An Empirical Model of Large-Batch Training. *arXiv preprint arXiv:1812.06162v1* (2018).
- [54] Hiroaki Mikami, Hisahiro Suganuma, Pongsakorn U.-Chupala, Yoshiki Tanaka, and Yuichi Kageyama. 2018. ImageNet/ResNet-50 Training in 224 Seconds. *CoRR abs/1811.05233* (2018). arXiv:1811.05233 <http://arxiv.org/abs/1811.05233>
- [55] T. Minka, J.M. Winn, J.P. Guiver, Y. Zaykov, D. Fabian, and J. Bronskill. 2018. /Infer.NET 0.3. Microsoft Research Cambridge. <http://dotnet.github.io/infer>.
- [56] Radford M Neal. 1993. *Probabilistic inference using Markov chain Monte Carlo methods*. Technical Report CRG-TR-93-1. Dept. of Computer Science, University of Toronto.
- [57] Radford M. Neal. 2011. MCMC using Hamiltonian dynamics. In *Handbook of Markov Chain Monte Carlo*, Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng (Eds.). Vol. 2. 2.
- [58] Feng Niu, Benjamin Recht, Christopher Re, and Stephen J. Wright. 2011. HOG-WILD!: A Lock-free Approach to Parallelizing Stochastic Gradient Descent. In *Proceedings of the 24th International Conference on Neural Information Processing Systems (NIPS'11)*. Curran Associates Inc., USA, 693–701. <http://dl.acm.org/citation.cfm?id=2986459.2986537>
- [59] X. Pan, J. Chen, R. Monga, S. Bengio, and R. Jozefowicz. 2017. Revisiting Distributed Synchronous SGD. *ArXiv e-prints* (Feb. 2017). arXiv:cs.DC/1702.05800

- [60] George Papamakarios and Iain Murray. 2016. Fast ϵ -free Inference of Simulation Models with Bayesian Conditional Density Estimation. In *Advances in Neural Information Processing Systems 29*, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Eds.). Curran Associates, Inc., 1028–1036.
- [61] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in PyTorch. In *NIPS 2017 Autodiff Workshop: The Future of Gradient-based Machine Learning Software and Techniques*, Long Beach, CA, US, December 9, 2017.
- [62] Michael E Peskin. 2018. *An introduction to quantum field theory*. CRC Press.
- [63] A Salam. 1968. Proceedings of the Eighth Nobel Symposium on Elementary Particle Theory, Relativistic Groups, and Analyticity, Stockholm, Sweden, 1968. (1968).
- [64] Christopher J. Shallue, Jaehoon Lee, Joseph Antognini, Jascha Sohl-Dickstein, Roy Frostig, and George E. Dahl. 2018. Measuring the Effects of Data Parallelism on Neural Network training. *arXiv preprint arXiv:1811.03600v2* (2018).
- [65] Samuel L Smith, Pieter-Jan Kindermans, Chris Ying, and Quoc V. Le. 2017. Don't Decay the Learning Rate, Increase the Batch Size. *arXiv preprint arXiv:1711.00489* (2017).
- [66] T. Smith, N. Maire, A. Ross, M. Penny, N. Chitnis, A. Schapira, A. Studer, B. Genton, C. Lengeler, F. Tediosi, and et al. 2008. Towards a comprehensive simulation model of malaria epidemiology and control. *Parasitology* 135, 13 (2008), 1507–1516. <https://doi.org/10.1017/S0031182008000371>
- [67] Sam Staton, Frank Wood, Hongseok Yang, Chris Heunen, and Ohad Kammar. 2016. Semantics for probabilistic programming: higher-order functions, continuous distributions, and soft constraints. In *2016 31st Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*. IEEE, 1–10.
- [68] John Sterman, Thomas Fiddaman, Travis Franck, Andrew Jones, Stephanie McCauley, Philip Rice, Elizabeth Sawin, and Lori Siegel. 2012. Climate interactive: the C-ROADS climate policy model. *System Dynamics Review* 28, 3 (2012), 295–305.
- [69] Michael Teng and Frank Wood. 2018. Bayesian Distributed Stochastic Gradient Descent. In *Advances in Neural Information Processing Systems*. 6380–6390.
- [70] Dustin Tran, Alp Kucukelbir, Adji B Dieng, Maja Rudolph, Dawen Liang, and David M Blei. 2016. Edward: A library for probabilistic modeling, inference, and criticism. *arXiv preprint arXiv:1610.09787* (2016).
- [71] Jan-Willem van de Meent, Brooks Paige, Hongseok Yang, and Frank Wood. 2018. An Introduction to Probabilistic Programming. *arXiv e-prints*, Article arXiv:1809.10756 (Sep 2018). [arXiv:stat.ML/1809.10756](https://arxiv.org/abs/1809.10756)
- [72] Martinus Veltman et al. 1972. Regularization and renormalization of gauge fields. *Nuclear Physics B* 44, 1 (1972), 189–213.
- [73] S Weinberg. 1967. *Phys. Rev. Lett* 19 (1967), 1264.
- [74] David Wingate, Andreas Stuhlmüller, and Noah Goodman. 2011. Lightweight implementations of probabilistic programming languages via transformational compilation. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*. 770–778.
- [75] Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google's neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144* (2016).
- [76] Yang You, Igor Gitman, and Boris Ginsburg. 2017. Large batch training of convolutional networks. *arXiv preprint arXiv:1708.03888* (2017).
- [77] S. Zheng, Q. Meng, T. Wang, W. Chen, N. Yu, Z.-M. Ma, and T.-Y. Liu. 2016. Asynchronous Stochastic Gradient Descent with Delay Compensation. *ArXiv e-prints* (Sept. 2016). [arXiv:cs.LG/1609.08326](https://arxiv.org/abs/1609.08326)

Appendix: Artifact Description/Artifact Evaluation

SUMMARY OF THE EXPERIMENTS REPORTED

For training results we ran the etalumis / pyprob code in a Docker container (links below) via Shifter on NERSC Cori and Edison machines (as well as other single node tests) as detailed in the paper text. `srun -nodes=$NODES -ntasks-per-node=2 shifter [exe]` With the following main arguments given below, and the number of nodes above; the optimizer; LR; Decay Type; BZPT and dataset locations below are all set depending on the experiment as indicated in the paper text. Inference and Dataset generation were performed using the same container code. Datasets can be generated using that code - but the primary dataset used for experiments in the paper is also made available at <https://portal.nersc.gov/project/dasrepo/etalumis-SC19/>

```
python /code/sherpa_probprog/scripts/sherpa_tau_decay.py
-m train -optimizer adam_larc -learning_rate_init ${LR}
-learning_rate_end ${LR_END} -learning_rate_scheduler
${LR_DECAY_TYPE} -distributed_num_buckets 1 -
distributed_backend mpi -address_dict /tmp/address_dict
-dataset_dir $DW_PERSISTENT_STRIPED_etalumis_traces_mar1/
etalumis_data_dec7_2018/traces_sorted_15M -dataset_valid_dir
$DW_PERSISTENT_STRIPED_etalumis_traces_mar1/
etalumis_data_dec7_2018/traces_sorted_valid -
network_dir mar19_1k_15M_networks_BZPT-
${BZPT}-Nodes-${SLURM_JOB_NUM_NODES}-LR-
${VAL}_converge_larc_${LR_DECAY_TYPE} -network_type
lstm -n $k -num_traces_end $m -log_file_name $LOG_NAME
-batch_size $BZPT -valid_every 4000000 -save_every 1200
```

ARTIFACT AVAILABILITY

Software Artifact Availability: Some author-created software artifacts are NOT maintained in a public repository or are NOT available under an OSI-approved license.

Hardware Artifact Availability: There are no author-created hardware artifacts.

Data Artifact Availability: Some author-created data artifacts are NOT maintained in a public repository or are NOT available under an OSI-approved license.

Proprietary Artifacts: None of the associated artifacts, author-created or otherwise, are proprietary.

List of URLs and/or DOIs where artifacts are available:

Base pyprob code (tagged version used in this paper)
↪ available at
<https://github.com/probprog/pyprob/releases/tag/0.13>
↪ .2.dev10

Container image with pyprob and sherpa-tau decay
↪ use-case
<https://hub.docker.com/r/etalumis/sherpa/tags/201904>
↪ 08

Container image used for experiments in this project
<https://hub.docker.com/r/wbhimji/etalumis-cori/tags/>
↪ v2019-03-20b
this contains pyprob and sherpa-tau-decay use-case
↪ together with the optimized pytorch-mpi used for
↪ NERSC experiments.

BASELINE EXPERIMENTAL SETUP, AND MODIFICATIONS MADE FOR THE PAPER

Relevant hardware details: CPUs: E5-2695 v2 @ 2.40GHz (12 cores/socket) E5-2698 v3 @ 2.30GHz (16 cores/socket) E5-2697A v4 @ 2.60GHz (16 cores/socket) Platinum 8170 @ 2.10GHz (26 cores/socket) Gold 6252 @ 2.10GHz (24 cores/socket); Systems: Cray XC30 Cray XC40; Filesystems Cray DataWarp; Sonnexion 2000 Lustre

Operating systems and versions: Ubuntu 18.04 (in container) ; Cray Linux Environment (base OS on NERSC machines)

Compilers and versions: Gcc-5 / g++-5 (in container); gcc (GCC) 7.3.0 20180125 (Cray Inc.) (Cori: used for optimized pytorch mpi)

Libraries and versions: Cray MPI :
/opt/cray/pe/mpt/7.7.3/gni/mpich-intel/16.0

Paper Modifications: All modifications are detailed in the paper and included in the containers given above.

Output from scripts that gathers execution environment information.

The environment used for running experiments was the
↪ standard NERSC environment on Cori and Edison as
↪ of Mar 20 - plus the container mentioned above
↪ (run in shifter). The NERSC environment is fairly
↪ verbose so the output of the
↪ collect_environment.sh script is not easily
↪ parseable but we provide it below as requested.

```
collect_environment.sh output (run on Cori compute
↪ node within the Docker/Shifter container used for
↪ this paper)
PE_MPICH_DEFAULT_DIR_CRAY_DEFAULT64=64
PE_PETSC_DEFAULT_VOLATILE_PrgENV=CRAY CRAY64 GNU
↪ GNU64 INTEL INTEL64
PE_NETCDF_DEFAULT_GENCOMPILERS_GNU=8.2 7.1 6.1 5.3 4.9
PE_TPSL_DEFAULT_GENCOMPILERS_INTEL_x86_skylake=16.0
PE_FFTW_DEFAULT_TARGET_x86_skylake=x86_skylake
```

```

PE_INTEL_DEFAULT_FIXED_PKGCONFIG_PATH=/opt/cray/pe/p
↪ arallel-netcdf/1.8.1.3/INTEL/16.0/lib/pkgconfig:
↪ /opt/cray/pe/netcdf-hdf5parallel/4.6.1.3/INTEL/1
↪ 6.0/lib/pkgconfig:/opt/cray/pe/netcdf/4.6.1.3/IN
↪ TEL/16.0/lib/pkgconfig:/opt/cray/pe/mpt/7.7.3/gn
↪ i/mpich-intel/16.0/lib/pkgconfig:/opt/cray/pe/hd
↪ f5-parallel/1.10.2.0/INTEL/16.0/lib/pkgconfig:/o
↪ pt/cray/pe/hdf5/1.10.2.0/INTEL/16.0/lib/pkgconfi
↪ g:/opt/cray/pe/ga/5.3.0.8/INTEL/18.0/lib/pkgconf
↪ ig
SLURM_STEP_NODELIST=nid00226
LIBRARY_PATH=/opt/intel/compilers_and_libraries_2018
↪ .1.163/linux/compiler/lib/intel64:/opt/intel/com
↪ pilers_and_libraries_2018.1.163/linux/mkl/lib/in
↪ tel64
PE_TPSSL_DEFAULT_GENCOMPILERS_CRAY_interlagos=8.6
NSG_HOME=/usr/syscom/nsg
PE_PETSC_DEFAULT_GENCOMPS_CRAY_sandybridge=86
HISTFILESIZE=100000
LESSOPEN=lessopen.sh %s
PE_FFTW_DEFAULT_TARGET_ivybridge=ivybridge
PE_TPSSL_DEFAULT_GENCOMPILERS_CRAY_x86_skylake=8.6
PE_CRAY_DEFAULT_FIXED_PKGCONFIG_PATH=/opt/cray/pe/pa
↪ rallel-netcdf/1.8.1.3/CRAY/8.6/lib/pkgconfig:/op
↪ t/cray/pe/netcdf-hdf5parallel/4.6.1.3/CRAY/8.6/l
↪ ib/pkgconfig:/opt/cray/pe/netcdf/4.6.1.3/CRAY/8.
↪ 6/lib/pkgconfig:/opt/cray/pe/hdf5-parallel/1.10.
↪ 2.0/CRAY/8.6/lib/pkgconfig:/opt/cray/pe/hdf5/1.1
↪ 0.2.0/CRAY/8.6/lib/pkgconfig:/opt/cray/pe/ga/5.3
↪ .0.8/CRAY/8.6/lib/pkgconfig
MINICOM=-c on
INPUTRC=/etc/inputrc
PE_TPSSL_64_DEFAULT_GENCOMPS_INTEL_haswell=160
RIVET_ANALYSIS_PATH=/code/rivetana
CRAY_UGNI_INCLUDE_OPTS=-I/opt/cray/ugni/6.0.14.0-6.0
↪ .7.0_23.1__gea11d3d.ari/include
ALLINEA_QUEUE_DLL=/opt/cray/pe/mpt/7.7.3/gni/mpich-i
↪ ntel/16.0/lib/libtvmppich.so.3.0.1
MAIL=/var/mail/USER
PE_TPSSL_64_DEFAULT_GENCOMPS_CRAY_haswell=86
PE_ENV=INTEL
MACHTYPE=x86_64-suse-linux
USER=USER
PE_MPICH_DEFAULT_GENCOMPILERS_CRAY=8.6
CRAY_PMI_INCLUDE_OPTS=-I/opt/cray/pe/pmi/5.0.14/incl
↪ ude
SSH_CLIENT=128.3.109.107 61346 22
SLURM_JOBID=20276515
PE_TPSSL_64_DEFAULT_GENCOMPS_GNU_x86_64=82 71 53 49
SLURM_CPU_BIND_LIST=0xFFFFFFFFFFFFFFFF
SLURM_JOB_USER=USER
SLURM_SPANK__SLURM_SPANK_OPTION_shifter_image=USER/e
↪ talumis-cori:v2019-03-20b
CVS_RSH=ssh
XTPE_NETWORK_TARGET=aries

CRAY_UDREG_POST_LINK_OPTS=-L/opt/cray/udreg/2.3.2-6.
↪ 0.7.0_33.18_g5196236.ari/lib64
PE_PETSC_DEFAULT_GENCOMPS_GNU_interlagos=71 53 49
PE_TPSSL_DEFAULT_GENCOMPILERS_INTEL_mic_knl=16.0
PE_MPICH_PKGCONFIG_LIBS=mpich
CRAY_LIBSCI_DIR=/opt/cray/pe/libsci/18.07.1
PE_FFTW_DEFAULT_TARGET_mic_knl=mic_knl
SLURM_PTY_PORT=38276
SRUN_DEBUG=3
HOSTNAME=cori12
BSCRATCH=/global/projectb/scratch/USER
LIBGL_DEBUG=quiet
SLURM_JOB_QOS=interactive
PE_PETSC_DEFAULT_GENCOMPILERS_INTEL_haswell=16.0
PE_TPSSL_DEFAULT_GENCOMPILERS_CRAY_mic_knl=8.6
PE_HDF5_PARALLEL_DEFAULT_GENCOMPILERS_GNU=8.2 7.1 6.1
↪ 5.3 4.9
LD_LIBRARY_PATH=/opt/udiImage/modules/mpich/lib64:/u
↪ sr/local/lib:/usr/local/nvidia/lib:/usr/local/n
↪ vidia/lib64
PKG_CONFIG_PATH_DEFAULT=/opt/cray/pe/papi/5.6.0.3/li
↪ b64/pkgconfig
PE_TPSSL_DEFAULT_GENCOMPS_GNU_haswell=82 71 53 49
PE_TPSSL_64_DEFAULT_GENCOMPILERS_INTEL_interlagos=16.0
PE_PAPI_DEFAULT_ACCEL_LIBS_nvidia35=-lcupti,-lcudar
↪ t,-lcuda
PE_TPSSL_DEFAULT_GENCOMPILERS_INTEL_sandybridge=16.0
SLURM_SRUN_COMM_PORT=37469
MPICH_ABORT_ON_ERROR=1
SLURM_JOB_NUM_NODES=1
LESSKEY=/etc/lesskey.bin
PE_PARALLEL_NETCDF_DEFAULT_VOLATILE_PKGCONFIG_PATH=/
↪ opt/cray/pe/parallel-netcdf/1.8.1.3/@PRGENV@/@PE
↪ _PARALLEL_NETCDF_DEFAULT_GENCOMPS@/lib/pkgconfig
SHLVL=4
PE_PETSC_DEFAULT_GENCOMPILERS_CRAY_haswell=8.6
PE_FFTW_DEFAULT_TARGET_sandybridge=sandybridge
PE_PETSC_DEFAULT_GENCOMPILERS_GNU_x86_64=7.1 5.3 4.9
XKEYSYMDB=/usr/X11R6/lib/X11/XKeysymDB
LESS=-M -I -R
PE_TPSSL_64_DEFAULT_GENCOMPILERS_CRAY_interlagos=8.6
SLURM_TASKS_PER_NODE=64
PE_HDF5_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe
↪ /hdf5/1.10.2.0/@PRGENV@/@PE_HDF5_DEFAULT_GENCOMP
↪ S@/lib/pkgconfig
CRAY_COOKIES=956956672,1409286144
PE_TRILINOS_DEFAULT_GENCOMPILERS_INTEL_x86_64=16.0
HOME=/global/homes/w/USER
XNLSPATH=/usr/share/X11/nls
PMI_CRAY_NO_SMP_ORDER=0
PE_LIBSCI_DEFAULT_OMP_REQUIRES_openmp=mp
RCLOCAL_PRGENV=true
INTEL_LICENSE_FILE=28518@crayintel.licenses.nersc.go
↪ v:28518@intel.licenses.nersc.gov
PE_TPSSL_DEFAULT_GENCOMPILERS_CRAY_sandybridge=8.6
PE_LIBSCI_GENCOMPS_GNU_x86_64=71 61 51 49

```

Etalumis: Bringing Probabilistic Programming to Scientific Simulators at Scale

```

PE_TPSL_64_DEFAULT_GENCOMPILERS_INTEL_x86_skylake=16
↪ .0
PE_PETSC_DEFAULT_GENCOMPS_GNU_mic_knl=53
PE_PETSC_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH:PE_LIBSC
↪ I:PE_HDF5_PARALLEL:PE_TPSL
PE_TRILINOS_DEFAULT_GENCOMPILERS_CRAY_x86_64=8.7
PE_PKGCONFIG_LIBS=mpich:AtpSigHandler:cray-rca:libsc
↪ i_mpi:libsci
PE_MPICH_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe/mpt/7.
↪ 7.3/gni/mpich-@PRGENV@PE_MPICH_DIR_DEFAULT64@/
↪ PE_MPICH_GENCOMPS@lib/pkgconfig
SLURM_TOPOLOGY_ADDR_PATTERN=switch.switch.node
SSH_TTY=/dev/pts/101
PE_SMA_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe/
↪ mpt/7.7.3/gni/sma@PE_SMA_DEFAULT_DIR_DEFAULT64@/
↪ lib64/pkgconfig
PE_TPSL_64_DEFAULT_GENCOMPILERS_CRAY_x86_skylake=8.6
PE_PETSC_DEFAULT_GENCOMPS_INTEL_x86_64=160
JRE_HOME=/usr/lib64/jvm/java/jre
CRAY_XPMEM_INCLUDE_OPTS=-I/opt/cray/xpmem/2.2.15-6.0
↪ .7.1_5.11__g7549d06.ari/include
PE_FFTW_DEFAULT_TARGET_abudhabi=abudhabi
SHIFTER_RUNTIME=1
MPICH_DIR=/opt/cray/pe/mpt/7.7.3/gni/mpich-intel/16.0
CRAY_PRE_COMPILE_OPTS=-hnetwork=aries
PE_PETSC_DEFAULT_GENCOMPILERS_INTEL_skylake=16.0
SLURM_CPU_BIND_TYPE=mask_cpu:
PE_NETCDF_DEFAULT_FIXED_PRGENV=CRAY_INTEL
PE_PETSC_DEFAULT_GENCOMPS_CRAY_x86_64=86
PE_LIBSCI_OMP_REQUIRES=
SLURM_PPIO_PROCESS=0
PE_PETSC_DEFAULT_GENCOMPILERS_CRAY_skylake=8.6
PE_LIBSCI_MODULE_NAME=cray-libsci/18.07.1
PE_PETSC_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/p
↪ e/petsc/3.8.4.0/complex/@PRGENV@PE_PETSC_DEFAU
↪ LT_GENCOMPS@PE_PETSC_DEFAULT_TARGET@lib/pkgco
↪ nfig
FPATH=/opt/cray/pe/modules/3.2.10.6/init/sh_funcs/n
↪ o_redirect:/opt/cray/pe/modules/3.2.10.6/init/sh
↪ _funcs/no_redirect:/opt/cray/pe/modules/3.2.10.6
↪ /init/sh_funcs/no_redirect
PE_PETSC_DEFAULT_GENCOMPS_GNU_sandybridge=71 53 49
PE_TPSL_DEFAULT_GENCOMPILERS_GNU_interlagos=8.2 7.1
↪ 5.3 4.9
PAGER=less
SLURM_SPANK_SHIFTER_IMAGE=7da07598b6f0c7229483c8dceb
↪ b766024910aa7fa6115485a1339c1d3ecb1165
SLURM_WORKING_CLUSTER=cori:ctl1:6817:8448
PE_TPSL_64_DEFAULT_GENCOMPILERS_INTEL_mic_knl=16.0
INTEL_VERSION=18.0.1.163
CRAYPE_DIR=/opt/cray/pe/craype/2.5.15
PE_TPSL_DEFAULT_GENCOMPILERS_GNU_x86_skylake=8.2 7.1
↪ 6.1
ATP_MRNET_COMM_PATH=/opt/cray/pe/atp/2.1.3/libexec/a
↪ tp_mrnet_commnnode_wrapper
PE_MPICH_NV_LIBS=

```

```

CRAY_ALPS_POST_LINK_OPTS=-L/opt/cray/alps/6.6.43-6.0
↪ .7.0_26.4__ga796da3.ari/lib64
MAN_POSIXLY_CORRECT=1
PE_TPSL_64_DEFAULT_GENCOMPILERS_CRAY_mic_knl=8.6
SHIFTER_IMAGE=7da07598b6f0c7229483c8dcebb766024910aa
↪ 7fa6115485a1339c1d3ecb1165
PE_LIBSCI_PKGCONFIG_VARIABLES=PE_LIBSCI_OMP_REQUIRES
↪ _@openmp@PE_SCI_EXT_LIBPATH:PE_SCI_EXT_LIBNAME
LIBSCI_BASE_DIR=/opt/cray/pe/libsci/18.07.1
G_BROKEN_FILENAMES=1
CSHRCREAD=true
PE_TPSL_DEFAULT_GENCOMPS_INTEL_interlagos=160
PE_TPSL_64_DEFAULT_GENCOMPILERS_INTEL_sandybridge=16
↪ .0
PE_SMA_DEFAULT_DIR_PGI_DEFAULT64=64
PE_TPSL_64_DEFAULT_GENCOMPS_GNU_haswell=82 71 53 49
SLURM_JOB_CPUS_PER_NODE=64
PE_NETCDF_HDF5PARALLEL_DEFAULT_GENCOMPILERS_GNU=8.2
↪ 7.1 6.1 5.3 4.9
SLURM_JOB_NAME=sh
NNTPSERVER=news
PE_TRILINOS_DEFAULT_VOLATILE_PRGENV=CRAY_GNU_INTEL
PE_MPICH_DEFAULT_GENCOMPILERS_GNU=7.1 5.1 4.9
PE_LIBSCI_DEFAULT_GENCOMPILERS_INTEL_x86_64=16.0
PE_MPICH_TARGET_VAR_nvidia20=-lcudart
_LMFILES=/opt/cray/pe/modulefiles/modules/3.2.10.6:
↪ /usr/syscom/nsg/modulefiles/nsg/1.2.0:/opt/modul
↪ efiles/intel/18.0.1.163:/opt/cray/pe/craype/2.5.
↪ 15/modulefiles/craype-network-aries:/opt/cray/pe
↪ /modulefiles/craype/2.5.15:/opt/cray/pe/modulefi
↪ les/cray-libsci/18.07.1:/opt/cray/ari/modulefile
↪ s/udreg/2.3.2-6.0.7.0_33.18__g5196236.ari:/opt/c
↪ ray/ari/modulefiles/ugni/6.0.14.0-6.0.7.0_23.1__
↪ gea11d3d.ari:/opt/cray/pe/modulefiles/pmi/5.0.14
↪ :/opt/cray/ari/modulefiles/dmapp/7.1.1-6.0.7.0_3
↪ 4.3__g5a674e0.ari:/opt/cray/ari/modulefiles/gni-
↪ headers/5.0.12.0-6.0.7.0_24.1__g3b1768f.ari:/opt
↪ /cray/ari/modulefiles/xpmem/2.2.15-6.0.7.1_5.11_
↪ _g7549d06.ari:/opt/cray/ari/modulefiles/job/2.2.
↪ 3-6.0.7.0_44.1__g6c4e934.ari:/opt/cray/ari/modul
↪ efiles/dvs/2.7_2.2.117-6.0.7.1_9.2__gf817677:/op
↪ t/cray/ari/modulefiles/alps/6.6.43-6.0.7.0_26.4_
↪ _ga796da3.ari:/opt/cray/ari/modulefiles/rca/2.2.
↪ 18-6.0.7.0_33.3__g2aa4f39.ari:/opt/cray/pe/modul
↪ efiles/atp/2.1.3:/opt/cray/pe/modulefiles/PrgEnv
↪ -intel/6.0.4:/opt/cray/pe/craype/2.5.15/modulefil
↪ es/craype-haswell:/opt/cray/pe/modulefiles/cray-
↪ mpich/7.7.3:/usr/common/software/modulefiles/alt
↪ d/2.0:/opt/modulefiles/Base-opts/2.4.135-6.0.7.0
↪ _38.1__g718f891.ari
PE_MPICH_DIR_CRAY_DEFAULT64=64
COLORTERM=1
SLURM_JOB_GID=68441
PE_NETCDF_DEFAULT_GENCOMPS_GNU=
PE_SMA_DEFAULT_DIR_CRAY_DEFAULT64=64
NERSC_HOST=cori

```

```

ALTD_VERBOSE=0
PE_TPSL_DEFAULT_GENCOMPS_CRAY_interlagos=86
ATP_POST_LINK_OPTS=-Wl,-L/opt/cray/pe/atp/2.1.3/libA
↪ pp/
PE_HDF5_PARALLEL_DEFAULT_FIXED_PRGENV=CRAY_INTEL
PE_INTEL_FIXED_PKGCONFIG_PATH=/opt/cray/pe/mpt/7.7.3
↪ /gni/mpich-intel/16.0/lib/pkgconfig
PE_MPICH_ALTERNATE_LIBS_multithreaded=_mt
CRAY_SITE_LIST_DIR=/etc/opt/cray/pe/modules
PE_TPSL_64_DEFAULT_GENCOMPILERS_CRAY_sandybridge=8.6
PE_GA_DEFAULT_GENCOMPILERS_GNU=5.3 4.9
CRAY_MPICH_DIR=/opt/cray/pe/mpt/7.7.3/gni/mpich-intel
↪ 1/16.0
PE_TPSL_DEFAULT_GENCOMPS_INTEL_x86_skylake=160
PE_LIBSCI_DEFAULT_GENCOMPILERS_CRAY_x86_64=8.6
CRAY_GNI_HEADERS_INCLUDE_OPTS=-I/opt/cray/gni-header
↪ s/5.0.12.0-6.0.7.0_24.1__g3b1768f.ari/include
PE_FFTW_DEFAULT_TARGET_x86_64=x86_64
SLURM_PROCID=0
PE_TPSL_DEFAULT_GENCOMPILERS_GNU_mic_knl=7.1 5.3
CRAY_UGNI_POST_LINK_OPTS=-L/opt/cray/ugni/6.0.14.0-6
↪ .0.7.0_23.1__gea11d3d.ari/lib64
INTEL_MAJOR_VERSION=18
PELOCAL_PRGENV=true
PE_TPSL_DEFAULT_GENCOMPILERS_INTEL_x86_64=16.0
SLURM_CPUS_ON_NODE=64
PE_FFTW_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH
CUDA_VERSION=9.2.148
ALTD_PATH=/usr/common/software/altd/2.0
PE_TPSL_DEFAULT_GENCOMPS_CRAY_x86_skylake=86
SLURM_JOB_ACCOUNT=m3224
PE_PKGCONFIG_PRODUCTS_DEFAULT=PE_PAPI
PE_PKGCONFIG_DEFAULT_PRODUCTS=PE_TRILINOS:PE_TPSL_64
↪ :PE_TPSL:PE_PETSC:PE_PARALLEL_NETCDF:PE_NETCDF_H
↪ DF5PARALLEL:PE_NETCDF:PE_MPICH:PE_LIBSCI:PE_HDF5
↪ _PARALLEL:PE_HDF5:PE_GA:PE_FFTW2:PE_FFTW
PE_NETCDF_DEFAULT_VOLATILE_PRGENV=GNU
PE_TPSL_DEFAULT_GENCOMPILERS_CRAY_x86_64=8.6
PE_MPICH_GENCOMPILERS_CRAY=8.6
PE_MPICH_FORTRAN_PKGCONFIG_LIBS=mpichf90
PE_MPICH_DEFAULT_GENCOMPS_CRAY=86
PE_PAPI_DEFAULT_PKGCONFIG_VARIABLES=PE_PAPI_ACCEL_LI
↪ BS_@accelerator@
CRAY_LIBSCI_PREFIX_DIR=/opt/cray/pe/libsci/18.07.1/I
↪ NTEL/16.0/x86_64
SLURM_STEP_LAUNCHER_PORT=37469
PE_PETSC_DEFAULT_GENCOMPILERS_GNU_haswell=7.1 5.3 4.9
CRAY_PMI_POST_LINK_OPTS=-L/opt/cray/pe/pmi/5.0.14/li
↪ b64
PE_TPSL_64_DEFAULT_GENCOMPILERS_GNU_interlagos=8.2
↪ 7.1 5.3 4.9
CRAY_PRGENVINTEL=loaded
TOOLMODULES=apprentice:apprentice2:atp:chapel:cray-1
↪ gdb:cray-snplauncher:craypat:craypkg-gen:ddt:gdb
↪ :iobuf:papi:perftools:perftools-lite:stat:totalv
↪ iew:xt-craypat:xt-lgdb:xt-papi:xt-totalview

```

```

PE_TPSL_DEFAULT_GENCOMPILERS_GNU_sandybridge=8.2 7.1
↪ 5.3 4.9
SCRATCH=/global/cscratch1/sd/USER
PE_FFTW_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe
↪ /fftw/3.3.8.1/@PE_FFTW_DEFAULT_TARGET@/lib/pkgco
↪ nfig
PE_TPSL_DEFAULT_GENCOMPS_INTEL_mic_knl=160
LOGNAME=USER
PE_LIBSCI_DEFAULT_VOLATILE_PRGENV=CRAY_GNU_INTEL
PE_TPSL_64_DEFAULT_GENCOMPILERS_GNU_x86_skylake=8.2
↪ 7.1 6.1
PE_TPSL_DEFAULT_GENCOMPS_CRAY_mic_knl=86
PE_HDF5_PARALLEL_DEFAULT_GENCOMPS_GNU=
SLURM_SPANK_SHIFTER_IMAGEREQUEST=USER/etalumis-cori:
↪ v2019-03-20b
PE_SMA_DEFAULT_COMPFLAG_GNU=-fcray-pointer
SLURM_SUBMIT_HOST=mom2
NVIDIA_REQUIRE_CUDA=cuda>=9.2
PE_PETSC_DEFAULT_GENCOMPS_INTEL_haswell=160
PE_TRILINOS_DEFAULT_GENCOMPILERS_GNU_x86_64=8.2 7.3
↪ 5.1 4.9
_=/collect_env.sh
PE_TPSL_64_DEFAULT_GENCOMPS_INTEL_interlagos=160
PE_PETSC_DEFAULT_GENCOMPS_CRAY_haswell=86
PE_PAPI_DEFAULT_ACCEL_FAMILY_LIBS_nvidia=-lcupti,-l
↪ cudart,-lcuda
PE_PETSC_DEFAULT_GENCOMPS_GNU_x86_64=71 53 49
DVS_INCLUDE_OPTS=-I/opt/cray/dvs/2.7.2.2.117-6.0.7.1
↪ _9.2__gf817677/include
PE_MPICH_TARGET_VAR_nvidia35=-lcudart
SHIFTER_IMAGEREQUEST=USER/etalumis-cori:v2019-03-20b
PE_TPSL_DEFAULT_GENCOMPS_INTEL_sandybridge=160
NVIDIA_DRIVER_CAPABILITIES=compute,utility
PKG_CONFIG_PATH=/opt/cray/rca/2.2.18-6.0.7.0_33.3__g
↪ 2aa4f39.ari/lib64/pkgconfig:/opt/cray/alps/6.6.4
↪ 3-6.0.7.0_26.4__ga796da3.ari/lib64/pkgconfig:/op
↪ t/cray/xpmem/2.2.15-6.0.7.1_5.11__g7549d06.ari/l
↪ ib64/pkgconfig:/opt/cray/gni-headers/5.0.12.0-6.
↪ 0.7.0_24.1__g3b1768f.ari/lib64/pkgconfig:/opt/cr
↪ ay/dmapp/7.1.1-6.0.7.0_34.3__g5a674e0.ari/lib64/
↪ pkgconfig:/opt/cray/pe/pmi/5.0.14/lib64/pkgconfi
↪ g:/opt/cray/ugni/6.0.14.0-6.0.7.0_23.1__gea11d3d
↪ .ari/lib64/pkgconfig:/opt/cray/udreg/2.3.2-6.0.7
↪ .0_33.18__g5196236.ari/lib64/pkgconfig:/opt/cray
↪ /pe/craype/2.5.15/pkg-config:/opt/cray/pe/iobuf/
↪ 2.0.8/lib/pkgconfig:/opt/cray/pe/fftw/2.1.5.9/li
↪ b/pkgconfig:/opt/cray/pe/atp/2.1.3/lib/pkgconfig
CPU=x86_64
PE_TPSL_64_DEFAULT_GENCOMPS_INTEL_x86_skylake=160
PE_PETSC_DEFAULT_GENCOMPILERS_GNU_skylake=6.1
XDG_SESSION_ID=45880
MORE=-s1
PE_FORTRAN_PKGCONFIG_LIBS=mpichf90
SLURM_NODELIST=nid00226
PE_LIBSCI_OMP_REQUIRES_openmp=_mp
PE_MPICH_MODULE_NAME=cray-mpich

```


Etalumis: Bringing Probabilistic Programming to Scientific Simulators at Scale

```
PE_TPSL_DEFAULT_GENCOMPS_CRAY_sandybridge=86
PE_HDF5_PARALLEL_DEFAULT_VOLATILE_PRGENV=GNU
PE_TRILINOS_DEFAULT_GENCOMPS_INTEL_x86_64=160
SLURM_STEP_RESV_PORTS=63098
PE_TPSL_64_DEFAULT_GENCOMPS_CRAY_interlagos=86
TERM=xterm-256color
PMI_NO_FORK=1
CRAY_LIBSCI_VERSION=18.07.1
PE_TPSL_DEFAULT_VOLATILE_PRGENV=CRAY CRAY64 GNU GNU64
↪ INTEL INTEL64
PE_TRILINOS_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH:PE_HD
↪ F5_PARALLEL:PE_NETCDF_HDF5PARALLEL:PE_LIBSCI:PE_
↪ TPSL
DMAPP_ABORT_ON_ERROR=1
LIBGL_ALWAYS_INDIRECT=1
PE_TPSL_64_DEFAULT_GENCOMPS_CRAY_x86_skylake=86
PE_MPICH_PKGCONFIG_VARIABLES=PE_MPICH_NV_LIBS_@accel
↪ erator@:PE_MPICH_ALTERNATE_LIBS_@multithreaded@:
↪ PE_MPICH_ALTERNATE_LIBS_@dpm@
QT_SYSTEM_DIR=/usr/share/desktop-data
PE_TPSL_64_DEFAULT_GENCOMPILERS_GNU_mic_knl=7.1 5.3
SLURM_NNODES=1
PE_SMA_DEFAULT_PKGCONFIG_VARIABLES=PE_SMA_COMPFLAG_@
↪ prgenv@
PE_NETCDF_HDF5PARALLEL_DEFAULT_FIXED_PRGENV=CRAY
↪ INTEL
PE_MPICH_NV_LIBS_nvidia20=-lcudart
ALT_LINKER=/usr/common/software/altd/2.0/bin/ld
PE_TPSL_64_DEFAULT_GENCOMPILERS_INTEL_x86_64=16.0
PE_TRILINOS_DEFAULT_GENCOMPS_CRAY_x86_64=87
SLURM_SPANK_SHIFTER_IMAGETYPE=id
CRAY_XPMEM_POST_LINK_OPTS=-L/opt/cray/xpmem/2.2.15-6
↪ .0.7.1_5.11_g7549d06.ari/lib64
PE_MPICH_DEFAULT_FIXED_PRGENV=INTEL
SLURM_JOB_ID=20276515
PE_TPSL_64_DEFAULT_GENCOMPILERS_CRAY_x86_64=8.6
PE_TRILINOS_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cra
↪ y/pe/trilinos/12.12.1.1/@PRGENV@/@PE_TRILINOS_DE
↪ FAULT_GENCOMPS@/@PE_TRILINOS_DEFAULT_TARGET@/lib
↪ /pkgconfig
HISTCONTROL=ignoredups:erasedups
LESS_ADVANCED_PREPROCESSOR=no
SLURM_CPU_BIND=quiet,mask_cpu:0xFFFFFFFFFFFFFFFF
SLURMD_NODENAME=nid00226
PE_PETSC_DEFAULT_GENCOMPS_INTEL_skylake=160
DVS_VERSION=0.9.0
KSH_AUTOLOAD=1
SLURM_JOB_MODEL_LIST=nid00226
PE_GA_DEFAULT_FIXED_PRGENV=CRAY INTEL
JAVA_BINDIR=/usr/lib64/jvm/java/bin
PE_PETSC_DEFAULT_GENCOMPILERS_INTEL_interlagos=16.0
PATH=/opt/conda/bin:/usr/local/nvidia/bin:/usr/local
↪ /cuda/bin:/usr/bin:/usr/local/sbin:/usr/local/bi
↪ n:/usr/sbin:/usr/bin:/sbin:/bin:/opt/udiImage/bin
PE_TPSL_64_DEFAULT_GENCOMPILERS_GNU_sandybridge=8.2
↪ 7.1 5.3 4.9
```

```
PRGENVMODULES=PrgEnv-cray:PrgEnv-gnu:PrgEnv-intel:Pr
↪ gEnv-pathscales:PrgEnv-pgi
PE_PETSC_DEFAULT_GENCOMPS_CRAY_skylake=86
PE_TPSL_DEFAULT_GENCOMPS_GNU_interlagos=82 71 53 49
JAVA_ROOT=/usr/lib64/jvm/java
PE_TPSL_64_DEFAULT_GENCOMPS_INTEL_mic_knl=160
SLURM_STEPID=0
PE_NETCDF_DEFAULT_REQUIRED_PRODUCTS=PE_HDF5
intel_already_loaded=0
PE_PARALLEL_NETCDF_DEFAULT_GENCOMPILERS_GNU=5.1 4.9
CSCRATCH=/global/cscratch1/sd/USER
PE_PETSC_DEFAULT_GENCOMPILERS_CRAY_interlagos=8.6
SLURM_GTIDS=0
PE_LIBSCI_DEFAULT_GENCOMPILERS_GNU_x86_64=7.1 6.1 5.1
↪ 4.9
PE_HDF5_DEFAULT_GENCOMPILERS_GNU=8.2 7.1 6.1 5.3 4.9
CRAY_MPICH_BASEDIR=/opt/cray/pe/mpt/7.7.3/gni
PE_TPSL_DEFAULT_GENCOMPILERS_INTEL_haswell=16.0
SLURM_SPANK_SHIFTER_SSH_PUBKEY=ssh-rsa
↪ AAAAB3NzaC1yc2EAAAADAQABAAQDonMQueK+xJzctc7Bq
↪ cq$+Bsr4I7+pfzrUcxXIONJf+b6IEVu+M23hLXse076DMrKz
↪ kPf1M9oL6RjDekjmEiIMBTwt+hbPrq1ULuwZTlUGc3MDgRB
↪ u/i6M3yFTBNC33dysrffJJaV6hTCuRRDFRnzik3HOJIUuwj/
↪ h0z1ZMtVNqi20Lx5aZlp340o7J/cS9Gr+B5v1K7Hf4Ggt0Rh
↪ JQeS8dqgjZEudML80RVpfAdrOc+Z2AXh87Ubg0s1PWsRBQ
↪ QEF2tMjdWnNRXsI72KPqa0LnboxdXUDUpibcl5MFoZ7AZWR4
↪ 1AahHEo4IGM7qv5BtxMD5oyahHdohv8z3g2f
↪ USER@nid00021

PE_TPSL_DEFAULT_GENCOMPS_GNU_x86_skylake=82 71 61
PE_FFTW_DEFAULT_TARGET_haswell=haswell1
CUDA_PKG_VERSION=9-2=9.2.148-1
XCURSOR_THEME=DMZ
PE_TPSL_64_DEFAULT_GENCOMPS_CRAY_mic_knl=86
ALTD_SELECT_ON=0
XDG_RUNTIME_DIR=/run/user/68441
CRAY_MPICH2_DIR=/opt/cray/pe/mpt/7.7.3/gni/mpich-int
↪ el/16.0
PE_PRODUCT_LIST=CRAYPE_HASWELL:CRAY_RCA:CRAY_ALPS:DV
↪ S:CRAY_XPMEM:CRAY_DMAPP:CRAY_PMI:CRAY_UGNI:CRAY_
↪ UDREG:CRAY_LIBSCI:CRAYPE:INTEL
PE_TPSL_DEFAULT_GENCOMPILERS_GNU_x86_64=8.2 7.1 5.3
↪ 4.9
SLURM_CHECKPOINT_IMAGE_DIR=/var/slurm/checkpoint
PE_TPSL_64_DEFAULT_GENCOMPS_INTEL_sandybridge=160
PE_NETCDF_HDF5PARALLEL_DEFAULT_GENCOMPS_GNU=
PE_LIBSCI_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH
PE_LIBSCI_GENCOMPILERS_INTEL_x86_64=16.0
PE_TPSL_DEFAULT_GENCOMPILERS_CRAY_haswell=8.6
ATP_HOME=/opt/cray/pe/atp/2.1.3
PE_NETCDF_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/
↪ pe/netcdf/4.6.1.3/@PRGENV@/@PE_NETCDF_DEFAULT_GE
↪ NCOMPS@/lib/pkgconfig
PE_MPICH_DEFAULT_GENCOMPS_GNU=71 51 49
PE_MPICH_GENCOMPILERS_GNU=7.1 5.1 4.9
PE_LIBSCI_DEFAULT_GENCOMPS_INTEL_x86_64=160
```

```

PE_MPICH_NV_LIBS_nvidia60=-lcudart
ALPS_LLI_STATUS_OFFSET=1
LANG=C.UTF-8
PE_LIBSCI_GENCOMPILERS_CRAY_x86_64=8.6
ZAP_LIBPATH=/opt/ovis/lib64/ovis-lib
ALTD_WORKDIR=/global/cscratch1/altd/logs
PE_FFTW2_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH
PE_GA_DEFAULT_GENCOMPS_GNU=53 49
PE_TPSL_64_DEFAULT_GENCOMPS_CRAY_sandybridge=86
LIBRARYMODULES=acml:alps:cray-dwarf:cray-fftw:cray-g
↳ a:cray-hdf5:cray-hdf5-parallel:cray-libsci:cray-
↳ libsci_acc:cray-mpich:cray-mpich-abi:cray-mpich2
↳ :cray-netcdf:cray-netcdf-hdf5parallel:cray-paral
↳ lel-netcdf:cray-petsc:cray-petsc-complex:cray-sh
↳ mem:cray-tpsl:cray-trilinos:cuda-toolkit:fftw:ga:
↳ hdf5:hdf5-parallel:iobuf:libfast:netcdf:netcdf-h
↳ df5parallel:ntk:onesided:papi:petsc:petsc-comple
↳ x:pmi:tpsl:trilinos:xt-libsci:xt-mpich2:xt-mpt:x
↳ t-papi
CRAY_MPICH_ROOTDIR=/opt/cray/pe/mpt/7.7.3
PE_LIBSCI_DEFAULT_GENCOMPS_CRAY_x86_64=86
PYTHONSTARTUP=/etc/pythonstart
PE_PETSC_DEFAULT_GENCOMPILERS_INTEL_mic_knl=16.0
PE_TPSL_64_DEFAULT_VOLATILE_PRGENV=CRAY CRAY64 GNU
↳ GNU64 INTEL INTEL64
HISTSIZE=100000
SLURM_STEP_NUM_NODES=1
CRAYPE_VERSION=2.5.15
LS_COLORS=no=00:fi=00:di=01;34:ln=00;36:pi=40;33:so=
↳ 01;35:do=01;35:bd=40;33;01:cd=40;33;01:or=41;33;
↳ 01:ex=00;32:*.cmd=00;32:*.exe=01;32:*.com=01;32:
↳ *.bat=01;32:*.btm=01;32:*.dll=01;32:*.tar=00;31:
↳ *.tbz=00;31:*.tgz=00;31:*.rpm=00;31:*.deb=00;31:
↳ *.arj=00;31:*.taz=00;31:*.lzh=00;31:*.lzma=00;31
↳ :*.zip=00;31:*.zoo=00;31:*.z=00;31:*.Z=00;31:*.g
↳ z=00;31:*.bz2=00;31:*.tb2=00;31:*.tz2=00;31:*.tb
↳ z2=00;31:*.xz=00;31:*.avi=01;35:*.bmp=01;35:*.fl
↳ i=01;35:*.gif=01;35:*.jpg=01;35:*.jpeg=01;35:*.m
↳ ng=01;35:*.mov=01;35:*.mpg=01;35:*.pcx=01;35:*.p
↳ bm=01;35:*.pgm=01;35:*.png=01;35:*.ppm=01;35:*.t
↳ ga=01;35:*.tif=01;35:*.xbm=01;35:*.xpm=01;35:*.d
↳ l=01;35:*.gl=01;35:*.wmv=01;35:*.aiff=00;32:*.au
↳ =00;32:*.mid=00;32:*.mp3=00;32:*.ogg=00;32:*.voc
↳ =00;32:*.wav=00;32:
PE_MPICH_NV_LIBS_nvidia35=-lcudart
CRAY_LD_LIBRARY_PATH=/opt/cray/pe/mpt/7.7.3/gni/mpi
↳ h-intel/16.0/lib:/opt/cray/rca/2.2.18-6.0.7.0_33
↳ .3_g2aa4f39.ari/lib64:/opt/cray/alps/6.6.43-6.0
↳ .7.0_26.4_ga796da3.ari/lib64:/opt/cray/xpmem/2.
↳ 2.15-6.0.7.1_5.11_g7549d06.ari/lib64:/opt/cray/
↳ dmapp/7.1.1-6.0.7.0_34.3_g5a674e0.ari/lib64:/op
↳ t/cray/pe/pmi/5.0.14/lib64:/opt/cray/ugni/6.0.14
↳ .0-6.0.7.0_23.1_ga11d3d.ari/lib64:/opt/cray/ud
↳ reg/2.3.2-6.0.7.0_33.18_g5196236.ari/lib64:/opt
↳ /cray/pe/libsci/18.07.1/INTEL/16.0/x86_64/lib
SLURM_CPU_BIND_VERBOSE=quiet

```

```

CRAY_DMAPP_INCLUDE_OPTS=-I/opt/cray/dmapp/7.1.1-6.0.
↳ 7.0_34.3_g5a674e0.ari/include
↳ -I/opt/cray/gni-headers/5.0.12.0-6.0.7.0_24.1_g
↳ 3b1768f.ari/include
CRAY_RCA_INCLUDE_OPTS=-I/opt/cray/rca/2.2.18-6.0.7.0
↳ _33.3_g2aa4f39.ari/include
↳ -I/opt/cray/krca/2.2.4-6.0.7.1_5.42_g8505b97.ar
↳ i/include
↳ -I/opt/cray/hss-devel/8.0.0/include
PE_PKGCONFIG_PRODUCTS=PE_MPICH:PE_LIBSCI
PE_PETSC_DEFAULT_GENCOMPILERS_CRAY_mic_knl=8.6
PE_TPSL_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH:PE_LIBSCI
PE_TPSL_DEFAULT_GENCOMPS_GNU_mic_knl=71 53
INTEL_MINOR_VERSION=18
PE_TPSL_DEFAULT_GENCOMPS_INTEL_x86_64=160
PE_LIBSCI_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/
↳ pe/libsci/18.07.1/@PRGENV@/PE_LIBSCI_DEFAULT_GE
↳ NCOMPS@/PE_LIBSCI_DEFAULT_TARGET@/lib/pkgconfig
SLURM_PTY_WIN_ROW=60
ALTD_ON=1
PE_NETCDF_HDF5PARALLEL_DEFAULT_VOLATILE_PRGENV=GNU
PE_MPICH_DEFAULT_VOLATILE_PRGENV=CRAY GNU
PE_HDF5_PARALLEL_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH
CUDNN_VERSION=7.2.1.38
PE_MPICH_GENCOMPS_CRAY=86
SLURM_CLUSTER_NAME=cori
PROFILEREAD=true
PE_TPSL_DEFAULT_GENCOMPS_CRAY_x86_64=86
SLURM_JOB_UID=68441
PE_PETSC_DEFAULT_GENCOMPS_GNU_haswell=71 53 49
PE_PAPI_DEFAULT_ACCELL_FAMILY_LIBS=
PE_PETSC_DEFAULT_GENCOMPILERS_INTEL_sandybridge=16.0
PE_GA_DEFAULT_VOLATILE_PRGENV=GNU
PE_PETSC_DEFAULT_GENCOMPILERS_CRAY_sandybridge=8.6
PE_TPSL_64_DEFAULT_GENCOMPS_GNU_interlagos=82 71 53
↳ 49
PE_TPSL_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe
↳ /tpsl/18.06.1/@PRGENV@/PE_TPSL_DEFAULT_GENCOMPS
↳ @/PE_TPSL_DEFAULT_TARGET@/lib/pkgconfig
PE_TPSL_DEFAULT_GENCOMPS_GNU_sandybridge=82 71 53 49
PE_PKGCONFIG_PATH=/opt/cray/pe/cti/1.0.7/lib/pkgcon
↳ fig:/opt/cray/pe/cti/1.0.6/lib/pkgconfig
SLURM_STEP_TASKS_PER_NODE=1
HOST=cori12
SLURM_UMASK=0002
LDMSD_PLUGIN_LIBPATH=/opt/ovis/lib64/ovis-ldms
SHELL=/bin/bash
PE_HDF5_PARALLEL_DEFAULT_VOLATILE_PKGCONFIG_PATH=/op
↳ t/cray/pe/hdf5-parallel/1.10.2.0/@PRGENV@/PE_HD
↳ F5_PARALLEL_DEFAULT_GENCOMPS@/lib/pkgconfig
SHIFTER_MODULE_MPICH=1
PE_TRILINOS_DEFAULT_GENCOMPS_GNU_x86_64=82 73 51 49
PE_TPSL_64_DEFAULT_GENCOMPS_GNU_x86_skylake=82 71 61
CSHEDIT=emacs
PE_LIBSCI_VOLATILE_PRGENV=CRAY GNU INTEL
PE_TPSL_64_DEFAULT_GENCOMPILERS_INTEL_haswell=16.0

```

Etalumis: Bringing Probabilistic Programming to Scientific Simulators at Scale

```

CXX=g++-5
PYTHON_VERSION=2.7
ALPS_APP_ID=20276515
SALLOC_IMMEDIATE=360
CRAY_NUM_COOKIES=2
LESSCLOSE=lessclose.sh %s %s
PE_MPICH_CXX_PKGCONFIG_LIBS=mpichcxx
PE_PAPI_DEFAULT_ACCEL_LIBS=
PE_PARALLEL_NETCDF_DEFAULT_FIXED_PRGENV=CRAY_INTEL
CRAY_CPU_TARGET=haswell
PE_TPSL_64_DEFAULT_GENCOMPILERS_GNU_x86_64=8.2 7.1
↳ 5.3 4.9
craype_already_loaded=0
MKLRROOT=/opt/intel/compilers_and_libraries_2018.1.16
↳ 3/linux/mkl
PE_TPSL_64_DEFAULT_GENCOMPILERS_CRAY_haswell=8.6
SLURM_LOCALID=0
SLURM_LAUNCH_NODE_IPADDR=10.128.5.222
MODULE_VERSION=3.2.10.6
CRAY_LIBSCI_BASE_DIR=/opt/cray/pe/libsci/18.07.1
CRAY_UDREG_INCLUDE_OPTS=-I/opt/cray/udreg/2.3.2-6.0.
↳ 7.0_33.18__g5196236.ari/include
PE_HDF5_DEFAULT_FIXED_PRGENV=CRAY_INTEL
SLURM_JOB_PARTITION=interactive
PE_PETSC_DEFAULT_GENCOMPILERS_GNU_interlagos=7.1 5.3
↳ 4.9
PE_PETSC_DEFAULT_GENCOMPS_GNU_skylake=61
SLURM_TASK_PID=65498
PE_TPSL_64_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH:PE_LIB
↳ SCI
WINDOWMANAGER=
PE_FFTW_DEFAULT_TARGET_broadwell=broadwell
PE_MPICH_FIXED_PRGENV=INTEL
PE_TPSL_64_DEFAULT_GENCOMPS_GNU_mic_knl=71 53
MODULE_VERSION_STACK=3.2.10.6
MPICH_MPIIO_DVS_MAXNODES=32
PE_TPSL_64_DEFAULT_GENCOMPS_INTEL_x86_64=160
ATP_IGNORE_SIGTERM=1
PE_TPSL_64_DEFAULT_GENCOMPS_CRAY_x86_64=86
PE_TPSL_DEFAULT_GENCOMPILERS_GNU_haswell=8.2 7.1 5.3
↳ 4.9
PE_NETCDF_HDF5PARALLEL_DEFAULT_REQUIRED_PRODUCTS=PE_
↳ HDF5_PARALLEL
OSTYPE=linux
SHMEM_ABORT_ON_ERROR=1
BASH_ENV=/global/homes/w/USER/.bashrc

```

```

USERMODULES=PrgEnv-cray:PrgEnv-gnu:PrgEnv-intel:PrgE
↳ nv-pathscale:PrgEnv-pgi:acml:alps:apprentice:app
↳ rentice2:atp:blcr:cce:chapel:cray-cddb:cray-fftw
↳ :cray-ga:cray-hdf5:cray-hdf5-parallel:cray-lgdb:
↳ cray-libsci:cray-libsci_acc:cray-mpich:cray-mpic
↳ h-compat:cray-mpich2:cray-netcdf:cray-netcdf-hdf
↳ 5parallel:cray-parallel-netcdf:cray-petsc:cray-p
↳ etsc-complex:cray-shmem:cray-snp-launcher:cray-tp
↳ sl:cray-trilinos:craypat:craype:craypkg-gen:cuda
↳ toolkit:ddt:fftw:ga:gcc:hdf5:hdf5-parallel:intel
↳ :iobuf:java:lgdb:libfast:libsci_acc:mpich1:netcd
↳ f:netcdf-hdf5parallel:netcdf-nofsync:netcdf-nofs
↳ ync-hdf5parallel:ntk:onesided:papi:parallel-netc
↳ df:pathscale:perftools:perftools-lite:petsc:pets
↳ c-complex:pgi:pmi:stat:totalview:tpsl:trilinos:x
↳ t-asyncpe:xt-craypat:xt-lgdb:xt-libsci:xt-mpich2
↳ :xt-mpt:xt-papi:xt-shmem:xt-totalview
CRAYPE_NETWORK_TARGET=aries
RCLOCAL_BASEOPTS=true
PE_PETSC_DEFAULT_GENCOMPS_INTEL_interlagos=160
LC_ALL=C.UTF-8
CRAY_IAA_INFO_FILE=/tmp/cray_iaa_info.20276515
PE_LIBSCI_DEFAULT_GENCOMPS_GNU_x86_64=71 61 51 49
GPG_TTY=/dev/pts/0
SLURM_SPANK_NERSC_ZONESORT_INTERVAL=0
PE_TPSL_64_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray
↳ /pe/tpsl/18.06.1/@PRGENV@64/@PE_TPSL_64_DEFAULT_
↳ GENCOMPS@/@PE_TPSL_64_DEFAULT_TARGET@/lib/pkgcon
↳ fig
JAVA_HOME=/usr/lib64/jvm/java
PE_LIBSCI_GENCOMPILERS_GNU_x86_64=7.1 6.1 5.1 4.9
PE_PETSC_DEFAULT_GENCOMPS_CRAY_interlagos=86
PWD=/global/homes/w/USER/etalumis_submitdir
PE_TPSL_64_DEFAULT_GENCOMPS_GNU_sandybridge=82 71 53
↳ 49
PE_LIBSCI_PKGCONFIG_LIBS=libsci_mpi:libsci
PE_CXX_PKGCONFIG_LIBS=mpichcxx
SLURM_TOPOLOGY_ADDR=s34.s10.nid00226
LOADED_MODULES=modules/3.2.10.6:nsg/1.2.0:intel/18.0.
↳ 1.163:craype-network-aries:craype/2.5.15:cray-li
↳ bsci/18.07.1:udreg/2.3.2-6.0.7.0_33.18__g5196236
↳ .ari:ugni/6.0.14.0-6.0.7.0_23.1__gea11d3d.ari:pm
↳ i/5.0.14:dmapp/7.1.1-6.0.7.0_34.3__g5a674e0.ari:
↳ gni-headers/5.0.12.0-6.0.7.0_24.1__g3b1768f.ari:
↳ xpmem/2.2.15-6.0.7.1_5.11__g7549d06.ari:job/2.2.
↳ 3-6.0.7.0_44.1__g6c4e934.ari:dvs/2.7.2.2.117-6.0
↳ .7.1_9.2__g817677:alps/6.6.43-6.0.7.0_26.4__ga7
↳ 96da3.ari:rca/2.2.18-6.0.7.0_33.3__g2aa4f39.ari:
↳ atp/2.1.3:PrgEnv-intel/6.0.4:craype-haswell:cray
↳ -mpich/7.7.3:altd/2.0:Base-opts/2.4.135-6.0.7.0_3
↳ 8.1__g718f891.ari
PE_PARALLEL_NETCDF_DEFAULT_GENCOMPS_GNU=51 49
ALTD_SELECT_USERS=
INTEL_PATH=/opt/intel/compilers_and_libraries_2018.1
↳ .163
SSH_CONNECTION=128.3.109.107 61346 128.55.209.23 22

```

```

PE_NETCDF_HDF5PARALLEL_DEFAULT_VOLATILE_PKGCONFIG_PA
↳ TH=/opt/cray/pe/netcdf-hdf5parallel/4.6.1.3/@PRG
↳ ENV@/@PE_NETCDF_HDF5PARALLEL_DEFAULT_GENCOMPS@/l
↳ ib/pkgconfig
PE_PETSC_DEFAULT_GENCOMPILERS_INTEL_x86_64=16.0
PE_MPICH_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/p
↳ e/mpt/7.7.3/gni/mpich-@PRGENV@/@PE_MPICH_DEFAULT_
↳ DIR_DEFAULT64@/@PE_MPICH_DEFAULT_GENCOMPS@/lib/p
↳ kgconfig
PKGCONFIG_ENABLED=1
XDG_DATA_DIRS=/usr/share
PE_HDF5_DEFAULT_GENCOMPS_GNU=
PE_PETSC_DEFAULT_GENCOMPILERS_GNU_mic_knl=5.3
PE_TPSSL_DEFAULT_GENCOMPS_INTEL_haswell=160
SLURM_MEM_PER_NODE=120832
XDG_CONFIG_DIRS=/etc/xdg
PE_FFTW_DEFAULT_TARGET_share=share
PE_PARALLEL_NETCDF_DEFAULT_VOLATILE_PRGENV=GNU
PMI_CONTROL_PORT=63098
PE_MPICH_GENCOMPS_GNU=71 51 49
TARGETMODULES=craype-abudhabi:craype-abudhabi-cu:cray
↳ ype-accel-host:craype-accel-nvidia20:craype-acce
↳ l-nvidia30:craype-accel-nvidia35:craype-barcelon
↳ a:craype-broadwell:craype-haswell:craype-hugepag
↳ es128K:craype-hugepages128M:craype-hugepages16M:
↳ craype-hugepages256M:craype-hugepages2M:craype-h
↳ ugepages32M:craype-hugepages4M:craype-hugepages5
↳ 12K:craype-hugepages512M:craype-hugepages64M:cray
↳ ype-hugepages8M:craype-intel-knc:craype-interlag
↳ os:craype-interlagos-cu:craype-istanbul:craype-i
↳ vybridge:craype-mc12:craype-mc8:craype-mic-knl:c
↳ raype-network-aries:craype-network-gemini:craype
↳ -network-infiniband:craype-network-none:craype-ne
↳ twork-seastar:craype-sandybridge:craype-shanghai
↳ :craype-target-compute_node:craype-target-local_
↳ host:craype-target-native:craype-xeon:xtpe-barce
↳ lona:xtpe-interlagos:xtpe-interlagos-cu:xtpe-ist
↳ anbul:xtpe-mc12:xtpe-mc8:xtpe-network-gemini:xtp
↳ e-network-seastar:xtpe-shanghai:xtpe-target-nati
↳ ve:xtpe-xeon
PE_LIBSCI_REQUIRED_PRODUCTS=PE_MPICH
PE_PETSC_DEFAULT_GENCOMPILERS_CRAY_x86_64=8.6
SDK_HOME=/usr/lib64/jvm/java
PE_LIBSCI_GENCOMPS_INTEL_x86_64=160
PE_TPSSL_DEFAULT_GENCOMPS_GNU_x86_64=82 71 53 49
PE_GA_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe/g
↳ a/5.3.0.8/@PRGENV@/@PE_GA_DEFAULT_GENCOMPS@/lib/
↳ pkgconfig
ALTD_SELECT_OFF_USERS=
PYTHONPATH=/opt/ovis/lib/python2.7/site-packages
PE_TPSSL_DEFAULT_GENCOMPS_CRAY_haswell=86
TZ=US/Pacific
PE_HDF5_DEFAULT_VOLATILE_PRGENV=GNU
NVIDIA_VISIBLE_DEVICES=all
PE_PETSC_DEFAULT_GENCOMPS_INTEL_mic_knl=160
CC=gcc-5

```

```

NCCL_VERSION=2.3.5
CPATH=/opt/intel/compilers_and_libraries_2018.1.163/
↳ linux/mkl/include
PE_LIBSCI_DEFAULT_OMP_REQUIRES=
PE_PETSC_DEFAULT_GENCOMPILERS_GNU_sandybridge=7.1
↳ 5.3 4.9
ALPS_APP_PE=0
PE_LIBSCI_GENCOMPS_CRAY_x86_64=86
HOSTTYPE=x86_64
SLURM_SRUN_COMM_HOST=10.128.5.222
CRAY_RCA_POST_LINK_OPTS=-L/opt/cray/rca/2.2.18-6.0.7
↳ .0_33.3__g2aa4f39.ari/lib64
↳ -lrca
PE_LIBSCI_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe/libsc
↳ i/18.07.1/@PRGENV@/@PE_LIBSCI_GENCOMPS@/@PE_LIBS
↳ CI_TARGET@/lib/pkgconfig
LIBSCI_VERSION=18.07.1
CRAY_ALPS_INCLUDE_OPTS=-I/opt/cray/alps/6.6.43-6.0.7
↳ .0_26.4__ga796da3.ari/include
PE_MPICH_ALTERNATE_LIBS_dpm=dpm
FROM_HEADER=
PE_MPICH_VOLATILE_PRGENV=CRAY GNU
LS_OPTIONS=-N --color=none -T 0
PE_LIBSCI_DEFAULT_PKGCONFIG_VARIABLES=PE_LIBSCI_DEFA
↳ ULT_OMP_REQUIRES_@openmp@:PE_SCI_EXT_LIBPATH:PE_
↳ SCI_EXT_LIBNAME
CRAY_DMAPP_POST_LINK_OPTS=-L/opt/cray/dmapp/7.1.1-6.
↳ 0.7.0_34.3__g5a674e0.ari/lib64
CRAY_MPICH2_VER=7.7.3
MANPATH=/usr/common/software/man:/usr/common/mss/man
↳ :/usr/common/nsg/man:/opt/cray/pe/mpt/7.7.3/gni/
↳ man/mpich:/opt/cray/pe/atp/2.1.3/man:/opt/cray/a
↳ lps/6.6.43-6.0.7.0_26.4__ga796da3.ari/man:/opt/c
↳ ray/job/2.2.3-6.0.7.0_44.1__g6c4e934.ari/man:/op
↳ t/cray/pe/pmi/5.0.14/man:/opt/cray/pe/libsci/18.
↳ 07.1/man:/opt/cray/pe/man/csm1version:/opt/cray/
↳ pe/craype/2.5.15/man:/opt/intel/compilers_and_li
↳ braries_2018.1.163/linux/man/common:/usr/syscom/
↳ nsg/man:/opt/cray/pe/modules/3.2.10.6/share/man:
↳ /usr/local/man:/usr/share/man:/opt/cray/share/ma
↳ n:/opt/cray/pe/man:/opt/cray/share/man:/opt/cray
↳ /share/man
MODULEPATH=/opt/cray/pe/craype/2.5.15/modulefiles:/o
↳ pt/cray/pe/modulefiles:/opt/cray/modulefiles:/op
↳ t/modulefiles:/usr/common/software/modulefiles:/
↳ usr/syscom/nsg/modulefiles:/usr/syscom/nsg/opt/m
↳ odulefiles:/usr/common/das/modulefiles:/usr/comm
↳ on/ftg/modulefiles:/opt/cray/craype/default/modu
↳ lefiles:/opt/cray/ari/modulefiles
PE_PETSC_DEFAULT_GENCOMPS_CRAY_mic_knl=86
SLURM_STEP_NUM_TASKS=1
SLURM_STEP_ID=0
SLURM_NODEID=0
SLURM_SUBMIT_DIR=/global/u1/w/USER/etalumis/code/she
↳ rpa/scripts/nersc
DVS_MAXNODES=1__

```

Etalumis: Bringing Probabilistic Programming to Scientific Simulators at Scale

```

PE_SMA_DEFAULT_COMPFLAG=
PE_PETSC_DEFAULT_GENCOMPS_INTEL_sandybridge=160
PE_TPSL_64_DEFAULT_GENCOMPILERS_GNU_haswell=8.2 7.1
↳ 5.3 4.9
JDK_HOME=/usr/lib64/jvm/java
PE_TPSL_DEFAULT_GENCOMPILERS_INTEL_interlagos=16.0
SLURM_PTY_WIN_COL=231
MODULESHOME=/opt/cray/pe/modules/3.2.10.6
PE_FFTW_DEFAULT_TARGET_interlagos=interlagos
Linux nid00226 4.4.103-6.38_4.0.153-cray_ari_c #1 SMP
↳ Thu Nov 1 16:05:05 UTC 2018 (6ef8fef) x86_64
↳ x86_64 x86_64 GNU/Linux
Architecture:      x86_64
CPU op-mode(s):    32-bit, 64-bit
Byte Order:        Little Endian
CPU(s):            64
On-line CPU(s) list: 0-63
Thread(s) per core: 2
Core(s) per socket: 16
Socket(s):         2
NUMA node(s):      2
Vendor ID:         GenuineIntel
CPU family:        6
Model:            63
Model name:        Intel(R) Xeon(R) CPU E5-2698 v3 @
↳ 2.30GHz
Stepping:          2
CPU MHz:           2301.000
CPU max MHz:       2301.0000
CPU min MHz:       1200.0000
BogoMIPS:          4600.29
Virtualization:    VT-x
L1d cache:         32K
L1i cache:         32K
L2 cache:          256K
L3 cache:          40960K
NUMA node0 CPU(s): 0-15,32-47
NUMA node1 CPU(s): 16-31,48-63
Flags:             fpu vme de pse tsc msr pae mce
↳ cx8 apic sep mtrr pge mca cmov pat pse36 clflush
↳ dts acpi mmx fxsr sse sse2 ss ht tm pbe syscall nx
↳ pdpe1gb rdtscp lm constant_tsc arch_perfmon pebs
↳ bts rep_good nopl xtopology nonstop_tsc
↳ aperfmperf eagerfpu pni pclmulqdq dtes64 monitor
↳ ds_cpl vmx smx est tm2 ssse3 sdbg fma cx16 xtpr
↳ pdcm pcid dca sse4_1 sse4_2 x2apic movbe popcnt
↳ tsc_deadline_timer aes xsave avx f16c rdrand
↳ lahf_lm abm ida arat epb invpcid_single pln pts
↳ dtherm spec_ctrl kaiser tpr_shadow vnmi
↳ flexpriority ept vpid fsgsbase tsc_adjust bmi1
↳ avx2 smep bmi2 erms invpcid cqm xsaveopt cqm_llc
↳ cqm_occup_llc
MemTotal:          131899768 kB
MemFree:           129215988 kB
MemAvailable:      128763396 kB
Buffers:           8796 kB

```

```

Cached:            592960 kB
SwapCached:        0 kB
Active:            243524 kB
Inactive:          393652 kB
Active(anon):      223016 kB
Inactive(anon):    198508 kB
Active(file):       20508 kB
Inactive(file):    195144 kB
Unevictable:       163316 kB
Mlocked:           8316 kB
SwapTotal:         0 kB
SwapFree:          0 kB
Dirty:             0 kB
Writeback:         0 kB
AnonPages:         198944 kB
Mapped:            65064 kB
Shmem:             230340 kB
Slab:              393268 kB
SReclaimable:      29344 kB
SUnreclaim:        363924 kB
KernelStack:       16912 kB
PageTables:         5352 kB
NFS_Unstable:      0 kB
Bounce:            0 kB
WritebackTmp:      0 kB
CommitLimit:       65949884 kB
Committed_AS:      700540 kB
VmallocTotal:      34359738367 kB
VmallocUsed:        0 kB
VmallocChunk:       0 kB
HardwareCorrupted: 0 kB
HugePages_Total:   0
HugePages_Free:    0
HugePages_Rsvd:    0
HugePages_Surp:    0
Hugepagesize:      2048 kB
DirectMap4k:        9204 kB
DirectMap2M:        1953792 kB
DirectMap1G:        134217728 kB
NAME      MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
loop0      7:0    0  192K  0 loop
loop1      7:1    0   4.5M  0 loop
loop2      7:2    0   2.4G  1 loop
loop3      7:3    0  23.8G  1 loop
loop4      7:4    0    4G   1 loop /opt/conda
loop5      7:5    0    1 loop
loop6      7:6    0    1 loop
loop7      7:7    0    1 loop
loop8      7:8    0    1 loop
loop9      7:9    0    1 loop
loop10     7:10   0    1 loop
loop11     7:11   0    1 loop
loop12     7:12   0    1 loop
loop13     7:13   0    1 loop
loop14     7:14   0    1 loop
loop15     7:15   0    1 loop
loop16     7:16   0    1 loop

```


loop17	7:17	0	1 loop	loop74	7:74	0	0 loop
loop18	7:18	0	1 loop	loop75	7:75	0	0 loop
loop19	7:19	0	1 loop	loop76	7:76	0	0 loop
loop20	7:20	0	1 loop	loop77	7:77	0	0 loop
loop21	7:21	0	1 loop	loop78	7:78	0	0 loop
loop22	7:22	0	1 loop	loop79	7:79	0	0 loop
loop23	7:23	0	1 loop	loop80	7:80	0	0 loop
loop24	7:24	0	1 loop	loop81	7:81	0	0 loop
loop25	7:25	0	1 loop	loop82	7:82	0	0 loop
loop26	7:26	0	1 loop	loop83	7:83	0	0 loop
loop27	7:27	0	1 loop	loop84	7:84	0	0 loop
loop28	7:28	0	1 loop	loop85	7:85	0	0 loop
loop29	7:29	0	1 loop	loop86	7:86	0	0 loop
loop30	7:30	0	1 loop	loop87	7:87	0	0 loop
loop31	7:31	0	1 loop	loop88	7:88	0	0 loop
loop32	7:32	0	0 loop	loop89	7:89	0	0 loop
loop33	7:33	0	0 loop	loop90	7:90	0	0 loop
loop34	7:34	0	0 loop	loop91	7:91	0	0 loop
loop35	7:35	0	0 loop	loop92	7:92	0	0 loop
loop36	7:36	0	0 loop	loop93	7:93	0	0 loop
loop37	7:37	0	0 loop	loop94	7:94	0	0 loop
loop38	7:38	0	0 loop	loop95	7:95	0	0 loop
loop39	7:39	0	0 loop	loop96	7:96	0	0 loop
loop40	7:40	0	0 loop	loop97	7:97	0	0 loop
loop41	7:41	0	0 loop	loop98	7:98	0	0 loop
loop42	7:42	0	0 loop	loop99	7:99	0	0 loop
loop43	7:43	0	0 loop	loop100	7:100	0	0 loop
loop44	7:44	0	0 loop	loop101	7:101	0	0 loop
loop45	7:45	0	0 loop	loop102	7:102	0	0 loop
loop46	7:46	0	0 loop	loop103	7:103	0	0 loop
loop47	7:47	0	0 loop	loop104	7:104	0	0 loop
loop48	7:48	0	0 loop	loop105	7:105	0	0 loop
loop49	7:49	0	0 loop	loop106	7:106	0	0 loop
loop50	7:50	0	0 loop	loop107	7:107	0	0 loop
loop51	7:51	0	0 loop	loop108	7:108	0	0 loop
loop52	7:52	0	0 loop	loop109	7:109	0	0 loop
loop53	7:53	0	0 loop	loop110	7:110	0	0 loop
loop54	7:54	0	0 loop	loop111	7:111	0	0 loop
loop55	7:55	0	0 loop	loop112	7:112	0	0 loop
loop56	7:56	0	0 loop	loop113	7:113	0	0 loop
loop57	7:57	0	0 loop	loop114	7:114	0	0 loop
loop58	7:58	0	0 loop	loop115	7:115	0	0 loop
loop59	7:59	0	0 loop	loop116	7:116	0	0 loop
loop60	7:60	0	0 loop	loop117	7:117	0	0 loop
loop61	7:61	0	0 loop	loop118	7:118	0	0 loop
loop62	7:62	0	0 loop	loop119	7:119	0	0 loop
loop63	7:63	0	0 loop	loop120	7:120	0	0 loop
loop64	7:64	0	0 loop	loop121	7:121	0	0 loop
loop65	7:65	0	0 loop	loop122	7:122	0	0 loop
loop66	7:66	0	0 loop	loop123	7:123	0	0 loop
loop67	7:67	0	0 loop	loop124	7:124	0	0 loop
loop68	7:68	0	0 loop	loop125	7:125	0	0 loop
loop69	7:69	0	0 loop	loop126	7:126	0	0 loop
loop70	7:70	0	0 loop	loop127	7:127	0	0 loop
loop71	7:71	0	0 loop				
loop72	7:72	0	0 loop				
loop73	7:73	0	0 loop				