Designing Neural Networks for Real-Time Systems

Pearce, Hammond; Yang, Xin; Roop, Partha S.; Katzef, Marc; Strøm, Tórir Biskopstø

Published in:
I E E E Embedded Systems Letters

Link to article, DOI:
[10.1109/LES.2020.3009910](https://doi.org/10.1109/LES.2020.3009910)

Publication date:
2021

Document Version
Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):
Pearce, H., Yang, X., Roop, P. S., Katzef, M., & Strøm, T. B. (2021). Designing Neural Networks for Real-Time Systems. *I E E E Embedded Systems Letters*, 13(3), 94-97. <https://doi.org/10.1109/LES.2020.3009910>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Designing Neural Networks for Real-Time Systems

Hammond Pearce, Xin Yang, Partha S. Roop, Marc Katzeff, Tórrur Biskopstø Strøm

Abstract—Artificial Neural Networks (ANNs) are increasingly being used within safety-critical Cyber-Physical Systems (CPSs). It is important to validate both the timing and functional correctness of these systems. However, most approaches in the literature consider guaranteeing only the functionality of ANN based controllers. This issue stems largely from the implementation strategies used within common neural network frameworks — their underlying source code is often simply unsuitable for formal techniques such as static timing analysis. As a result, developers of safety-critical CPS must rely on informal techniques such as measurement based approaches to prove correctness, techniques that provide weak guarantees at best.

In this work we address this challenge. We propose a design pipeline whereby neural networks trained using the popular deep learning framework Keras are compiled to functionally equivalent C code. This C code is restricted to simple constructs that may be analysed by existing static timing analysis tools. As a result, if compiled to a suitable time-predictable platform all execution bounds may be statically derived.

To demonstrate the benefits of our approach we execute an ANN trained to drive an autonomous vehicle around a race track. We compile the ANN to the Patmos time-predictable controller, and show that we can derive worst case execution timings.

I. INTRODUCTION

In safety-critical Cyber-Physical Systems (CPSs), timing correctness can be as important as functional correctness. Consider the case published in [1] where an autonomous *F1/10* vehicle is using LiDAR to drive safely through a racetrack, as depicted in Figure 1. A corner approaches. The vehicle must react appropriately, ensuring that two properties are met: firstly, that the corner is detected and the steering is changed to avoid crashing; and secondly, that the decision to steer the car around the corner is completed in a timely fashion. In other words, if the car controller takes too long to process the change in the road, the output of that controller is *incorrect* and will potentially lead to a crash.

Figure 1: Autonomous driving course for *F1/10* car

H. Pearce, X. Yang, P. S. Roop, and M. Katzeff were with the Department of Electrical, Computer, and Software Engineering, University of Auckland, Auckland, New Zealand.

T. B. Strøm was with the Technical University of Denmark, Denmark.

Corresponding Email: hammond.pearce@auckland.ac.nz

Manuscript received April 01, 2020; revised June 30, 2020.

Unfortunately, while traditional control code for CPS such as autonomous vehicles is designed to be amenable to both functional verification (e.g. ensure that the car will turn) and timing verification (e.g. ensure that this decision will take place quickly enough), often modern control systems involving Artificial Intelligence (AI) do not. Instead, popular learning frameworks such as Keras [2] and Caffe [3] rely on training/validation, simulation, and deployment testing, approaches which only provide weak guarantees of correctness [4]. In addition, approaches such as those provided through TensorFlow Lite [5] are popularising the usage of these kinds of neural networks in embedded devices without further examination of the safety implications involved.

To address this there is a push in the literature to bring formal methodologies into the AI domain [6], [7], especially within the usage of Artificial Neural Networks (ANNs) such as Multi-Layer Perceptrons (MLPs) and Convolutional Neural Networks (CNNs). For example, the case study in [1] was verified using reachability analysis, whereby the sigmoid activation function inside the neurons of the MLP controller were converted into a non-linear hybrid system before being analysed using a tool they called *Flow** [8], [9]. Despite this, the issue of timing verification of ANN-based controllers has received scant attention [4].

In this work we seek to address this challenge by providing a new design pipeline for converting neural networks trained in Keras to time-predictable C code. This generated code, once deployed upon a suitable time-predictable architecture, may then be statically analysed to formally derive timing bounds.

The rest of this paper is organised as follows. Section II discusses the state of the art in this area. Section III discusses our solution to this challenge, where Keras neural networks are compiled to functionally equivalent C code. Finally, Section IV presents our evaluation and Section V concludes.

II. BACKGROUND

While it is sometimes common to consider *real-time* execution as simply *fast* execution, for safety-critical applications such as those found within CPSs it is not enough to simply execute your software quickly [10], as rarely-executed branches not covered by testing may hide the presence of delays or other *timing anomalies* [11]. As such, for a system to be truly *real-time*, it must be proven to meet its timing requirements via techniques such as Static Timing Analysis (STA). Unfortunately, while there has been plenty of work in the literature focussed on increasing the execution speed of neural networks (for instance the work presented in [12], where Redmon et al. present a fast execution framework for CNNs), the issue of formal timing predictability for ANNs has received less attention.

In many cases, this is related to the underlying implementations of the neural network libraries used. For STA of software to take place, two key requirements must be met [10]. Firstly, it must be possible to convert program binaries into traversable control flow graphs. This can be very difficult given the presence of program interpreters (for instance if the neural network is implemented in a language such as Python); the presence of a complex runtime and/or Operating System (OS); and/or the presence of dynamic control flows (for instance code generalised to run many neural networks rather than being specialised to one). Secondly, the code must be targeted at a time-predictable processor so that execution times can be derived for the paths through the control flow graph.

As such, software implementations of neural networks can be very difficult to time. For instance, TensorFlow Lite [5], which is one of the most popular tools used by industry to embed neural networks in cyber-physical devices, features a runtime and requires Linux to operate — and as such it is not amenable to formal STA.

To address this challenge, in [4] Roop et al. proposed synchronous execution of neural networks, where they may be designed with synchronous programming languages such as Esterel and then executed on a time-predictable processor such as Patmos [13]. However, their approach was limited, and not all networks could be timed.

Other approaches, observing that neural networks can be compiled to hardware rather than software (such as in Had-doc2 [14]) focus on timing hardware implementations of neural networks [15]. Other hardware-based approaches also exist, although these mainly focus on performance [16] rather than predictability. In addition, deploying neural networks to reconfigurable hardware is much more complex process than that provided through software tools such as Keras.

III. KERAS2C: OUR DESIGN APPROACH

There is no fundamental reason that implementations of neural networks must be so difficult to time. At their heart, neural networks such as MLPs and CNNs are just different kinds of feed-forward mathematical functions. In this section we detail how a neural network, trained in Keras [2] to drive the autonomous vehicle track in Figure 1, may be formally implemented in a time-predictable manner.

A. Design Process

The Keras2C design process is depicted in Figure 2. Users will (1) design and train their networks using the Keras

Figure 2: Keras2C Design Process

framework then (2) use the normal validation process to check functionality. From (3) the saved Keras network files (an architecture.json which describes the network structure, and a weights.h5 storing the network's trained weights) are given to the Keras2C tool which generates a folder of .c and .h files. These may then be (5) compiled with the open-source Patmos-clang compiler, and the output of the compilation further analysed with the existing STA tool *platin* [17] to derive Worst-Case Execution Time (WCET). If satisfactory, these files may then be considered the final implementation (6) of the neural network.

B. Compiling Multi-Layer Perceptrons (MLPs)

Let us examine this more formally for an MLP type neural network, as this is suitable for driving the autonomous vehicle around the track in Figure 1. An MLP can be encoded as Definition III.1.

Definition III.1. An MLP can be formalised as a tuple $M = \langle I, O, L, N, \alpha, B, C, W, f \rangle$, where:

- I is a finite collection of n input variables with its domain being $\mathbf{I} = \mathbb{R}^n$
- O is a finite collection of m output variables with its domain being $\mathbf{O} = \mathbb{R}^m$
- L is a set of layers where l_0 represents the input layer and $l_{|L|-1}$ represents the output layer.
- N is a set of neurons.
- $\alpha : L \rightarrow N$ is a neuron mapping function $\alpha(l_n)$ where no neuron can appear in more than one layer. Neurons in the input layer are mapped to inputs I (i.e. $|\alpha(l_0)| = |I|$), and neurons in the output layer are mapped to outputs O (i.e. $|\alpha(l_{|L|-1})| = |O|$).
- $B : N \rightarrow \mathbb{R}$ is a function $B(n_i)$ which returns a real-valued bias for a given neuron.
- $C \subseteq N \times N$ is an ordered set of inter-layer connections $c_{i,j}$ between neurons n_i and n_j such that $n_i \in \alpha(l_k)$ and $n_j \in \alpha(l_{k+1})$ (i.e. connections can only go from a neuron in one layer to a neuron in the next layer).
- $W : C \rightarrow \mathbb{R}$ is a function $W(c_{i,j})$ which returns a real-valued weighting for a given connection $c_{i,j} \in C$.
- $f : \mathbb{R} \rightarrow \mathbb{R}$ is the neuron activation function (e.g. ReLU).

An MLP according to Definition III.1 can be straightforwardly executed using Algorithm 1. Here, lines 1-3 set the input neuron layer to simply be their input values from I . Lines 4-12 then compute the value of all other neurons (in the hidden and output layers), by first setting all other neurons to their initial value i.e. their bias (line 6) and then summing their source neuron values multiplied by their connection weights (lines 7-9) before passing the sum into their activation function (line 10). Finally, the output collection O can be produced (lines 13-15).

It is straightforward to see that if Algorithm 1 is directly realised as static C code with bounded for loops then this execution will be time-predictable. However, while this approach is fundamentally the same as any used by the popular machine learning frameworks such as Keras, they focus on dynamic approaches for improving average-case execution

Algorithm 1 Execution of an MLP M

```

1: for all  $n_i \in \alpha(l_0)$  do
2:    $n_i \leftarrow \text{from\_input}(I, n_i)$ 
3: end for
4: for each  $l \in L \setminus l_0$  do
5:   for each  $n_i \in \alpha(l)$  do
6:      $n_i \leftarrow B(n_i)$ 
7:     for each  $c_{h,i} \in C$  from  $n_h$  to  $n_i$  do
8:        $n_i \leftarrow n_i + n_h * W(c_{h,i})$ 
9:     end for
10:     $n_i \leftarrow f(n_i)$ 
11:   end for
12: end for
13: for all  $n_i \in \alpha(l_k)$  where  $k = |L| - 1$  do
14:    $\text{set\_output}(O, n_i)$ 
15: end for

```

time, featuring optimisations that can introduce unexpected delays and timing anomalies, and effectively prevent STA.

In our case, we instead use the methodology in Algorithm 1 to generate functionally equivalent C code from the saved Keras network files. The C code is static, only capable of executing the ANN it is generated from, and no operating system or runtimes are required.

C. Other network types

Definition III.1 and Algorithm 1 are designed for static execution of MLPs. However, any stateless feedforward neural networks (such as CNNs) may be executed in a similar manner, as they too may be simply considered sequences of mathematical operations.

IV. EVALUATION

In this section, we evaluate the efficacy of our approach for a range of available benchmarks against two other methodologies. Firstly, we examine our approach in a formal setting by benchmarking it against the existing results in [4] for time-predictable execution of neural networks on the Patmos architecture. Then, we examine the raw scalability and raw performance of our approach more generally by benchmarking it against the popular TensorFlow Lite framework.

A. Comparison with Synchronous ANNs (SANNs)

[4] introduces a synchronous framework for execution of ANNs that they term SANN. In this section we will benchmark our approach against their proposal.

1) *Methodology*: Using the SANN approach, neural networks are trained offline and then compiled in a framework involving the Esterel programming language to C code. However, while they could execute both CNNs and MLPs, only MLPs were compiled to predictable C code, which could then be executed over the time-predictable Patmos architecture, detailed here.

2) *The Patmos time-predictable processor*: The Patmos [13] processor is one part of the larger T-CREST [18] project, which features a series of time-predictable hardware and associated tools, including interconnect, memories,

and software tool-chain including the LLVM-based compiler Patmos-clang and the WCET analysis tool Platin [19].

Patmos is a RISC style processor optimised for WCET analysis. An example of this optimization is the use of special cache types to aid WCET analysis: a stack cache, reserved for stack allocated data, and a method cache for full function caching. If the function does not fit into a cache block, it is broken into smaller functions by the compiler.

3) *Benchmarking*: For this approach we compare four benchmarks from SANN [4]. Of the four, only XOR and Adder had formal analysis presented in the original paper, however our tool can generate predictable code for any MLP and CNN from Keras. In addition, we also present the results for F1/10 running on Patmos, presented both with and without the fix16 library used in the SANN benchmarks. All benchmarks were analysed over a single-core 50MHz Patmos processor.

Benchmark	Number of connections	SANN WCET (ms) from [4]	Keras2C WCET (ms)
XOR	9	0.82	0.1
Adder	15	0.49	0.1
Rabbit	576	Unavailable	2.93
Wolf	840	Unavailable	4.03
F1/10 (fix16)	24,320	N/A	82.83
F1/10 (float)	24,320	N/A	3,235.39

Table I: Benchmarking Keras2C against SANN. Presented times are “per invocation”.

As can be seen in Table I, the Keras2C code runs faster than the SANN approach. We believe this is due to overheads introduced by the Esterel compilation process. Our approach is also more amenable to STA, as the CNN library utilised in the original approach is not suitable for analysis.

4) *Results for the F1/10 case study*: Consider the track depicted in Figure 1. It is 20m by 10m with a width of 1.5m. The racing car starts from the mid-point of the left side of the track, drives forwards and makes two left turns, and will finally stop at the mid-point of the right side of the track.

The vehicle is initially stationary, and it accelerates until its velocity reaches 2.4ms^{-1} . It typically takes the car approximately 20s to complete the track.

In the original case study, the car’s controller is expected to emit a control output every 100ms (i.e. every 0.24m of travel distance at top speed). If this deadline is not met then an error will be introduced. If the deadline continues to fail to be met then further errors are introduced. Eventually these errors may build up sufficiently that the car will crash.

Utilising our approach, the F1/10 case study is guaranteed to execute in 82.83ms when utilising the fix16 library for fixed point math. However, if this library is not used, timing analysis becomes inaccurate due to the difficulties involved with predicting execution time for floating point computations [20], and meeting the 100ms deadline cannot be guaranteed.

5) *Analysing the LeNet-5 benchmark*: To demonstrate our ability to analyse CNNs, we also analyse the 60,000 parameter benchmark LeNet-5 [21]. Given the same Patmos core at 50MHz, the toolchain derives a WCET of 751.84ms (when using fix16) and 30,381ms (when using floats).

B. Performance Comparison with TensorFlow Lite

TensorFlow Lite is a popular tool for executing neural networks on platforms such as smartphones and embedded computers (e.g. the ARM-based Raspberry Pi). It uses a dynamic runtime with a background garbage collector to execute its neural networks, and requires a full underlying OS. As such it is not amenable to STA, unlike Keras2C.

Benchmark	Number of connections	TensorFlow Lite		Keras2C	
		Avg. (ms)	WCET (ms)	Avg. (ms)	WCET (ms)
XOR	9	0.027	15.10	0.0023	0.06
Adder	15	0.024	10.06	0.0026	0.064
Rabbit	576	0.046	10.15	0.026	0.14
Wolf	840	0.048	13.07	0.037	0.13
F1/10	24320	0.096	2.269	0.98	1.401

Table II: Benchmarking Keras2C against TensorFlow Lite. Presented times are “per invocation”. Lower is better.

Table II presents our comparison with TensorFlow Lite on a Raspberry Pi Model 3B, where we measured the execution times of the network over a random input using the system clock. This is repeated for one million invocations. For fairness, we omit time taken for the I/O of both frameworks as well as the first 100 ticks of TensorFlow Lite due to its significant “wind up” time. As can be seen, while our approach is faster in the average case for the smaller networks, TensorFlow Lite outperforms our C code for the larger F1/10 network. This is likely due to its internal optimisation routines which focus on average case performance. However, these optimisations come at a cost, with unexpected delays in the worst case and overheads introduced for small networks. As such, in all benchmarks our framework (as it is optimised for timing predictability) had a lower measured WCET.

V. CONCLUSIONS

Neural networks are increasingly being adopted within CPS. However, while these systems might have safety-critical timing requirements, there are few approaches for verifying timing properties of neural network based code. Through a series of benchmarks we have demonstrated our approach, Keras2C. It addresses the challenge through a simple-to-use toolchain for converting neural networks trained in Keras to C code amenable to static timing analysis. Keras2C is made freely available at <https://aitransformer.com>.

Future work could examine how certain optimisations could be integrated to improve average-case performance without compromising our timing predictability.

REFERENCES

- [1] R. Ivanov, T. J. Carpenter, J. Weimer, R. Alur, G. J. Pappas, and I. Lee, “Case study: Verifying the safety of an autonomous racing car with a neural network controller,” *arXiv preprint arXiv:1910.11309*, 2019.
- [2] F. Chollet. (2015) Keras. [Online]. Available: <https://github.com/fchollet/keras>
- [3] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell, “Caffe: Convolutional architecture for fast feature embedding,” in *Proceedings of the 22Nd ACM International Conference on Multimedia*, ser. MM ’14. New York, NY, USA: ACM, 2014, pp. 675–678. [Online]. Available: <http://doi.acm.org/10.1145/2647868.2654889>

- [4] P. S. Roop, H. Pearce, and K. Monadjem, “Synchronous neural networks for cyber-physical systems,” in *2018 16th ACM/IEEE International Conference on Formal Methods and Models for System Design (MEMOCODE)*, 10 2018, pp. 1–10.
- [5] TensorFlow. (2019, 09) TensorFlow: Lite. [Online]. Available: <https://www.tensorflow.org/lite>
- [6] S. Tripakis, “Data-driven and model-based design,” in *2018 IEEE Industrial Cyber-Physical Systems (ICPS)*. IEEE, 2018, pp. 103–108.
- [7] G. Schwalbe and M. Schels, “A Survey on Methods for the Safety Assurance of Machine Learning Based Systems,” in *10th European Congress on Embedded Real Time Software and Systems (ERTS 2020)*, Toulouse, France, Jan. 2020. [Online]. Available: <https://hal.archives-ouvertes.fr/hal-02442819>
- [8] R. Ivanov, J. Weimer, R. Alur, G. J. Pappas, and I. Lee, “Verisig: verifying safety properties of hybrid systems with neural network controllers,” *CoRR*, vol. abs/1811.01828, 2018. [Online]. Available: <http://arxiv.org/abs/1811.01828>
- [9] X. Chen, E. Ábrahám, and S. Sankaranarayanan, “Flow*: An analyzer for non-linear hybrid systems,” in *Computer Aided Verification*, N. Sharygina and H. Veith, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 258–263.
- [10] A. Roychoudhury, *Embedded systems and software validation*. Morgan Kaufmann, 2009.
- [11] J. Reineke, B. Wachter, S. Thesing, R. Wilhelm, I. Polian, J. Eisinger, and B. Becker, “A Definition and Classification of Timing Anomalies,” in *6th International Workshop on Worst-Case Execution Time Analysis (WCET’06)*, ser. OpenAccess Series in Informatics (OASICS), F. Mueller, Ed., vol. 4. Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2006. [Online]. Available: <http://drops.dagstuhl.de/opus/volltexte/2006/671>
- [12] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, “You only look once: Unified, real-time object detection,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [13] M. Schoeberl, W. Puffitsch, S. Hepp, B. Huber, and D. Prokesch, “Patmos: A time-predictable microprocessor,” *Real-Time Systems*, vol. 54(2), pp. 389–423, Apr 2018.
- [14] K. Abdelouahab, M. Pelcat, J. Srot, C. Bourrasset, and F. Berry, “Tactics to directly map CNN graphs on embedded FPGAs,” *IEEE Embedded Systems Letters*, vol. 9, no. 4, pp. 113–116, Dec 2017.
- [15] N. Allen, Y. Raje, J. W. Ro, and P. Roop, “A compositional approach for real-time machine learning,” in *Proceedings of the 17th ACM-IEEE International Conference on Formal Methods and Models for System Design*, ser. MEMOCODE ’19. New York, NY, USA: Association for Computing Machinery, 2019. [Online]. Available: <https://doi.org/10.1145/3359986.3361204>
- [16] V. Sze, Y. Chen, T. Yang, and J. S. Emer, “Efficient processing of deep neural networks: A tutorial and survey,” *Proceedings of the IEEE*, vol. 105, no. 12, pp. 2295–2329, 2017.
- [17] P. Puschner, D. Prokesch, B. Huber, J. Knoop, S. Hepp, and G. Gebhard, “The T-CREST approach of compiler and WCET-analysis integration,” in *16th IEEE International Symposium on Object/component/service-oriented Real-time distributed Computing (ISORC 2013)*, 06 2013, pp. 1–8.
- [18] M. Schoeberl, S. Abbaspour, B. Akesson, N. Audsley, R. Capasso, J. Garside, K. Goossens, S. Goossens, S. Hansen, R. Heckmann, S. Hepp, B. Huber, A. Jordan, E. Kasapaki, J. Knoop, Y. Li, D. Prokesch, W. Puffitsch, P. Puschner, A. Rocha, C. Silva, J. Sparsø, and A. Tocchi, “T-CREST: Time-predictable multi-core architecture for embedded systems,” *Journal of Systems Architecture*, vol. 61, no. 9, pp. 449–471, 2015.
- [19] S. Hepp, B. Huber, J. Knoop, D. Prokesch, and P. P. Puschner, “The platin tool kit - the T-CREST approach for compiler and WCET integration,” in *Proceedings 18th Kolloquium Programmiersprachen und Grundlagen der Programmierung, KPS 2015, Pörtlach, Austria, October 5-7, 2015*, 2015.
- [20] M. Andryscio, A. Nötzli, F. Brown, R. Jhala, and D. Stefan, “Towards verified, constant-time floating point operations,” in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS ’18. New York, NY, USA: ACM, 2018, pp. 1369–1382. [Online]. Available: <http://doi.acm.org/10.1145/3243734.3243766>
- [21] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, “Gradient-based learning applied to document recognition,” *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.