MNRL and MNCaRT

An Open-Source, Multi-Architecture State Machine Research and Execution Ecosystem

Kevin Angstadt*, Jack Wadden†, Vinh Dang†, Ted Xie†, Dan Kramp†,
Westley Weimer*, Mircea Stan†, Kevin Skadron†

*Computer Science and Engineering
University of Michigan, Ann Arbor, MI 48109
{angstadt, weimerw}@umich.edu

†Department of Computer Science
University of Virginia, Charlottesville, VA 22904
{wadden, vqd8a, ted.xie, dankramp, mircea, skadron}@virginia.edu

***Abstract*—We present MNCaRT, a comprehensive software ecosystem for the study and use of automata processing across hardware platforms. Tool support includes manipulation of automata, execution of complex machines, high-speed processing of NFAs and DFAs, and compilation of regular expressions. We provide engines to execute automata on CPUs (with VASim and Intel Hyperscan), GPUs (with custom DFA and NFA engines), and FPGAs (with an HDL translator). We also introduce MNRL, an open-source, general-purpose and extensible state machine representation language developed to support MNCaRT. The representation is flexible enough to support traditional finite automata (NFAs, DFAs) while also supporting more complex machines, such as those which propagate multi-bit signals between processing elements. We hope that our ecosystem and representation language stimulates new efforts to develop efficient and specialized automata processing applications.**

I. Introduction

Years of research and development have resulted in high-throughput automata processing architectures and software engines [1]–[3]. This has lead to the discovery of non-obvious use-cases and application domains for finite automata, such as natural language processing [4], network security [5], graph analytics [6], high-energy physics [7], bioinformatics [8]–[10], pseudo-random number generation and simulation [11], data-mining [12], [13], and machine learning [14].

Unfortunately, the software frameworks for the construction, manipulation, and translation of automata are frustratingly fractured (e.g. have inconsistent serialization formats) and restrictively licensed (e.g., Micron licenses a comprehensive SDK, but it is closed-source and specifically targets their D480 Automata Processor, or AP [2]). While these tools are useful for developing applications for the AP, the tools do not allow researchers to easily evaluate designs across hardware platforms, such as CPUs, GPUs, and FPGAs. The tools also cannot be easily extended to support new architectures and automata paradigms. Instead, a general and extensible framework is needed to enable the development of platform-independent applications and to support experimental automata designs.

We present MNCaRT (the MNRL Network Computation and Research Testbed, pronounced "minecart") [15], a suite of tools for creating, manipulating, and executing automata.[1] MNCaRT collects a diverse set of automata processing tools and algorithms into a central location and will grow as new tools are developed. We currently provide support for compiling state machines from Perl compatible regular expressions (PCRE) [16] to MNRL, high-speed execution of NFAs and DFAs using Intel Hyperscan [3], and optimization and simulation of experimental automata designs with the Virtual Automata Simulator (VASim) [17]. Further, we provide back-ends for executing DFAs on GPUs (DFAGE) [18], FPGAs [19], and exploring routing constraints for experimental spatial architectures via the Automata-to-Routing (ATR) tool [20]. All tools in MNCaRT are publicly available (typically under BSD licenses), and available pre-installed in a Linux container, allowing both academics and industry experts to contribute to, and use, the ecosystem.

To support our ecosystem, we have created MNRL, the MNRL Network Representation Language (pronounced "mineral"), a JSON-based, open-source language to support the development of, and experimentation with, new automata-based applications and architectures. MNRL allows a user to define a *network* (or collection) of MNRL *nodes*, which represent the states within automata. Each node stores configuration information (such as node type, name, etc.) and connections to other nodes within the network. The language specification is general, allowing state machines other than finite automata to be represented. We provide initial definitions for traditional finite automata states, homogeneous states, up-counters, and Boolean logic in the MNRL specification; additional node types may be defined by the user for specific applications.

To summarize, this works presents the following:

- MNCaRT, an comprehensive repository of compatible tools for development, visualization, and analysis of automata processing on CPUs, GPUs, and FPGAs.
- MNRL, an extensible, open-source JSON specification for representing state machines.
- Extensions to Intel's Hyperscan PCRE engine, supporting compilation to and execution of MNRL files.
- Updated versions of VASim, REAPR, DFAGE, and iNFAnt2, which support reading and writing of MNRL files.

II. Background and Related Work

A finite automaton includes of a set of states and a set of transitions defining how the states become active based

[1] https://github.com/kevinaangstadt/mncart

on symbols observed in an input stream [21]. In a non-deterministic finite automaton (NFA), it is possible to transition to multiple states on the same input symbol. Automata are often represented as a graph, defining the topological layout of the computation. Computation is therefore decoupled from the definition of the state machine, allowing for a common execution engine to process arbitrary automata, improving code reuse and reducing sources for bugs. Automata can also be represented as a set of regular expressions, which define the search pattern the automata recognize.

In the remainder of this section, we briefly highlight some existing automata processing engines and discuss limitations of current automata representation languages.

Automata Processing Engines. Micron's D480 AP [2] is a custom hardware accelerator which directly executes homogeneous finite automata.[2] Becchi et al. have developed a set of tools and algorithms for efficient CPU-based automata processing [22], and other CPU engines include Intel's Hyperscan [3] and Google's RE2 [23]. Automata processing engines have also been developed for GPUs and FPGAs (e.g., [24], [25], and [26]). Unfortunately, existing engines do not share a common automata representation, making cross-architecture comparison and development of automata-based algorithms challenging and time-consuming.

Limitations of Automata Representation Languages. The Automata Network Markup Language (ANML) is a proprietary description language developed for the Micron D480 AP. Licensing restrictions make the language challenging to use for prototyping new automata elements, and additional annotations cannot be added to elements in ANML while maintaining support for current tools. Therefore it is not a good choice for unifying automata processing engines.

Becchi et al.'s tools use a simple NFA representation based on the theoretic definition of NFAs and cannot be easily extended to support more complex state machines. The language is custom, and there is no support in general-purpose programming languages for reading and manipulating these files.

Regular expressions are commonly used to generate automata, but are difficult to develop and maintain. Many applications (e.g., particle tracking, motif searchers, and rule mining) would be represented by non-intuitive regular expressions that are often exhaustive enumerations of all possible matches. Additionally, programming of regular expressions can be extremely error-prone due to variations in regular expression syntax, which leads to high rates of runtime exceptions [27].

While other automata representation languages exist (e.g., Dot and JFLAP), these present similar licensing, generalizability, and maintainability challenges.

III. MNRL: A JSON-Based Automata Language

We have developed MNRL, an extensible, open-source automata representation language, which allows for the topological specification of a collection of finite state machines using JSON syntax. While JSON is supported by most common general-purpose programming languages, we provide C++ and Python bindings to support additional validation checks.

```
1  {
2    "id": "0t_15l_5r",
3    "type": "hState",
4    "enable": "onActivateIn",
5    "report": true,
6    "inputDefs": [
7      {
8        "width": 1,
9        "portId": "i"
10     }
11   ],
12   "outputDefs": [
13     {
14       "width": 1,
15       "activate": [],
16       "portId": "o"
17     }
18   ],
19
20   "attributes": {
21     "reportId": 5,
22     "latched": false,
23     "symbolSet": "[\\xFF]"
24   }
25 }
```

Fig. 1. Sample MNRL homogeneous hState Node. The node is enabled (performs computation) only after an incoming edge is active (line 4), and this node matches against the input character \xFF (line 23). When this occurs, the node generates a report signal (line 5). Lines 6-11 define a single input port for incoming edges. Lines 12-18 define a single output port for outgoing edges. The array on line 15 is empty, indicating that there are no outgoing edges.

It is important to note that the MNRL format specifies the layout of a machine but does not specify how elements behave, allowing many types of state machines to be represented, including traditional NFAs [21] and homogeneous NFAs [28].[3] Behavior is left for the execution engine to specify and implement (allowing MNRL to be an extremely flexible file format). Therefore, MNRL is similar in intent to the Unified Modeling Language (UML), in which developers describe and design software systems while eliding implementation details [29].

A. MNRL Format

A MNRL file contains a single MNRL *network*—a collection of one or more state machines that are executed in parallel using the same input. The file contains an array of MNRL nodes, which define each element in the network. An example node is given in Figure 1. A node consists of:

- A unique identifier
- A node type (state, homogeneous state, up counter, boolean, etc.)
- How the node is enabled (e.g. on the start input processing or when an incoming edge is active)

[2]In a homogeneous NFA, all incoming transitions to any given state *must* occur on the same input character.

[3]MNRL is general enough to represent more powerful machines (e.g. push-down automata, cellular automata, and Turing machines).

- Whether the node reports (generates an output signal) when activated
- An array of input ports, each with a unique ID and specified width (number of wires)
- An array of output ports, each with a unique ID, specified width, and list of connected nodes
- Custom attributes, specific to each element type

A developer can encode the topological layout of the state machines within the network and to specify the sort of behavior the underlying execution engine should assign to each node. The implementation of behavior is *not* defined in the MNRL file; instead, the computation engine that processes a MNRL network is responsible for specifying the semantics for each node type. Therefore, node types and execution engines are typically co-designed. If an engine needs information (e.g. symbol sets for matching against an input stream) to process a node, this configuration can be embedded in a MNRL node's attributes. For the standard node types, we have specified additional attributes to support their respective expected behaviors.[4]

B. Extending the MNRL Schema

MNRL is designed to be extensible, enabling research on new, custom automata functionality and allows researchers to quickly define custom attributes for new node types. We provide the specification of MNRL as a JSON schema [31], which allows for validation of file syntax. Because custom node types become part of the JSON schema, prototype extensions to the MNRL format can still be statically checked with minimal effort from the developer. The MNRL file format could easily be extended to support additional node types such as non-deterministic counters [32], and stacks (to support push-down automata). Because MNRL supports variable-width ports, it is also possible to represent elements that share more than a single bit of data with elements downstream.

IV. The MNCaRT Ecosystem

Our goal with this work is to enable the development of a rich, vibrant ecosystem of compatible tools for manipulating and executing automata. We are collecting these tools in an umbrella repository, the MNRL Network Computation and Research Testbed (or MNCaRT). By keeping tools catalogued in a single location, we hope to maintain the interoperability of tools and reduce fracturing in the ecosystem. We also provide a Linux container configured to use all of the MNCaRT tools.[5]

Figure 2 describes the interaction between tools provided with MNCaRT. Our ecosystem supports workflows beginning with high-level languages, such as PCRE, and ending with execution on CPUs, GPUs, and FPGAs. We also support execution on Micron's Automata Processor via conversion to Micron's Automata Network Markup Language (ANML). Additionally, we provide compatible benchmarks for testing and experimenting with tools in MNCaRT. In this section, we briefly describe the tools that make up the initial release of MNCaRT.

[4]For additional details, please see Angstadt et al. [30].

[5]https://hub.docker.com/r/kevinaangstadt/mncart

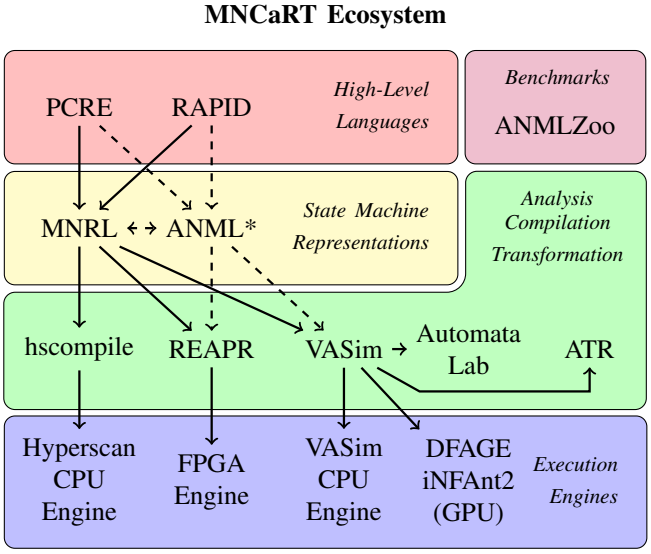

Fig. 2. Tools supplied as part of MNCaRT. These fall into four categories: front-end representations (both high-level and representation languages), benchmarks, transformation and compilation tools, and hardware and software execution engines.

A. High-Level Languages

Regular Expressions. Our framework supports programming models that represent pattern searches at higher levels of abstraction. We compile PCRE to MNRL files using Intel Hyperscan's parsing and compilation routines [3]. Hyperscan is an open-source, industry-standard regular expression processing library supported by Intel. The tool returns a graph representation of the compiled state machine, which we traverse to generate a MNRL file.

RAPID. RAPID is a high-level programming language for execution of sequential pattern-matching applications [33]. This C-like language is extended with three keywords to support parallel matching of patterns against a single data stream as well as sliding window pattern recognition. We have extended the RAPID compiler to emit MNRL files, allowing for high-level programming within the MNCaRT ecosystem.

B. Benchmarks

ANMLZoo. The ANMLZoo benchmark suite contains a diverse set of automata applications and associated input stimuli [18]. Applications range from configurable, synthetic benchmarks to algorithms that are not easily represented by regular expressions and can therefore have vastly different execution characteristics. We have generated MNRL representations for all benchmarks in the suite.

C. State Machine Representations

MNRL. Our JSON-based automata representation language provides the glue for MNCaRT ecosystem. This allows the various tools to function seamlessly with each other. Further details regarding MNRL may be found in Section III.

ANML. MNCaRT also supports Micron's Automata Network Markup Language (ANML) and therefore allows performance comparisons with the D480 AP. We provide this support via a translation to MNRL. This translation is supported for MNRL node types provided with our default distribution of the language. Because element types in ANML are dictated by the design of the AP, our translator does not support custom node types.

D. Analysis, Transformation, and Compilation

Hyperscan Compilation. We provide an extension to Hyperscan (**hscompile**) that parses MNRL files and compiles finite automata to a serialized Hyperscan pattern database, allowing offline compilation. Hyperscan is an industry standard automata processing toolchain that offers a state-of-the-art, high-performance CPU automata processing engine.

VASim. We have extended VASim [17] to support parsing of MNRL files. VASim is a general-purpose framework for automata simulation, optimization, transformation, and performance modeling. The tool enables prototyping, debugging, simulation, and analysis of automata-based applications and architectures. Additionally, VASim can parse Micron ANML files, allowing for conversion with MNRL.

Automata Lab. Automata Lab is a web-based graphical environment for visualizing, editing, and simulating finite automata [34]. The tool uses VASim to manipulate automata, and the resulting state machines are displayed graphically, allowing for user interaction. Users may upload MNRL files or choose from applications in the ANMLZoo benchmark suite.

REAPR. REAPR [19] is a design automation tool for generating highly-efficient FPGA-based automata kernels. REAPR generates RTL representations of homogeneous finite automata, a class of automata also implemented by Micron's AP. Additionally, REAPR generates a reporting architecture, which allow for pattern matches discovered by the executing automata to be communicated back to the host system.

Automata-to-Routing. ATR utilizes the Versatile Place and Route (VPR) tool to model spatial automata processing architectures [35]. We use VASim to emit VPR-readable circuits of MNRL networks and provide guidance to construct custom, parameterizable, spatial architecture description files to accept these custom state machine circuits. ATR is thus capable of modeling spatial architectures that are purpose-built to accept MNRL state machines.

E. Execution Engines

Hyperscan CPU Engine. We provide a tool (**hsrun**) for processing MNRL files against an input stream using the Hyperscan execution core. This tool deserializes the Hyperscan pattern database and node mapping produced by **hscompile**. The tool then scans the given input file against the database and prints out human-readable reporting information (e.g. MNRL ID and input stream offset). If multiple compiled MNRL files and/or input files are passed to **hsrun**, the tool will execute all pairings of the files using a supplied number of threads.

VASim CPU Engine. In addition to support for transformation and analysis of finite automata, VASim supports simulation of a diverse set of finite automata models. While Hyperscan achieves higher throughput, VASim's modular design allows for quick prototyping to test new automata elements and designs, such as those including custom compute units.

FPGA Engine. In addition to generating hardware NFA kernels, REAPR can also generate a full platform execution environment for certain automata applications. The REAPR platform has been demonstrated to offer up to $183\times$ speedup over best-effort CPU implementations [19].

GPU Engines (DFAGE and iNFAnt2). MNCaRT contains both a GPU-based DFA engine (DFAGE) and NFA engine (iNFAnt2). The NFA engine was described previously by Wadden et al. [18]; we therefore briefly describe DFAGE. Use of DFAGE first requires automata to be transformed into one or more DFAs using VASim. Note that this process is performed offline by the CPU. Each DFA consists of a state transition table and an acceptance vector, both of which are stored in the GPU's global memory. A transition table is represented by a 2-D array containing the next state identifiers for every pair of current state identifier and input symbol. Similar to previous implementations, our DFA matching engine supports multi-packet processing to take advantage of the extreme parallelism of GPU architectures.

V. Conclusions

We present MNCaRT, a suite of tools for analyzing, executing, and transforming automata networks. We support execution of MNRL networks on CPUs, GPUs, and FPGAs, and we provide a workflow for execution on Micron's AP. Support for high-level pattern-matching languages, such as PCRE and RAPID is also provided as part of MNCaRT. Finally, we allow for design space exploration through analysis functionality in the VASim and ATR tools. We hope that this suite of tools will enable innovation and in this emerging, and important application domain.

Acknowledgment

This work was supported in part by grants from the NSF (CCF-1116673, CCF-1629450, CCF-1619123, CNS-1619098), AFRL (FA8750-15-2-0075), Jefferson Scholars Foundation, Achievement Rewards for College Scientists

(ARCS) Foundation, a grant from Xilinx, and support from C-FAR, one of six centers of STARnet, a Semiconductor Research Corporation program sponsored by MARCO and DARPA. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of AFRL.

REFERENCES

[1] Titan IC Systems, "Helios RXPF soft IP for FPGA security analytics acceleration," http://titan-ic.com/products/helios-rxpf, 2017.

[2] P. Dlugosch, D. Brown, P. Glendenning, M. Leventhal, and H. Noyes, "An efficient and scalable semiconductor architecture for parallel automata processing," *IEEE Transactions on Parallel and Distributed Systems*, vol. 25, no. 12, pp. 3088–3098, 2014. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6719386

[3] Intel, "Hyperscan," https://01.org/hyperscan, 2017.

[4] K. Zhou, J. J. Fox, K. Wang, D. E. Brown, and K. Skadron, "Brill tagging on the Micron Automata Processor," in *Proceedings of the 9th IEEE International Conference on Semantic Computing*, 2015, pp. 236–239. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7050812

[5] I. Roy, A. Srivastava, M. Nourian, M. Becchi, and S. Aluru, "High performance pattern matching using the automata processor," in *Proceedings of the IEEE International Parallel and Distributed Processing Symposium*, ser. IPDPS '16, 2016, pp. 1123–1132.

[6] I. Roy, N. Jammula, and S. Aluru, "Algorithmic techniques for solving graph problems on the Automata Processor," in *Proceedings of the IEEE International Parallel and Distributed Processing Symposium*, ser. IPDPS '16, May 2016, pp. 283–292.

[7] M. H. Wang, G. Cancelo, C. Green, D. Guo, K. Wang, and T. Zmuda, "Using the Automata Processor for fast pattern recognition in high energy physics experiments - a proof of concept," *Nuclear Instruments and Methods in Physics Research*, 2016. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0168900216306891

[8] I. Roy and S. Aluru, "Finding motifs in biological sequences using the Micron Automata Processor," in *Proceedings of the 28th IEEE International Parallel and Distributed Processing Symposium*, 2014, pp. 415–424. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6877275

[9] I. Roy, "Algorithmic techniques for the micron automata processor," Ph.D. dissertation, Georgia Institute of Technology, 2015.

[10] T. Tracy II, M. Stan, N. Brunelle, J. Wadden, K. Wang, K. Skadron, and G. Robins, "Nondeterministic finite automata in hardware—the case of the Levenshtein automaton," *Architectures and Systems for Big Data (ASBD), in conjunction with ISCA*, 2015.

[11] J. Wadden, N. Brunelle, K. Wang, M. El-Hadedy, G. Robins, M. Stan, and K. Skadron, "Generating efficient and high-quality pseudo-random behavior on automata processors," in *2016 IEEE 34th International Conference on Computer Design (ICCD)*, Oct 2016, pp. 622–629.

[12] K. Wang, M. Stan, and K. Skadron, "Association rule mining with the Micron Automata Processor," in *Proceedings of the 29th IEEE International Parallel & Distributed Processing Symposium*, 2015. [Online]. Available: http://www.cap.virginia.edu/sites/cap.virginia.edu/files/kwang_arm_submitted.pdf

[13] K. Wang, E. Sadredini, and K. Skadron, "Sequential pattern mining with the Micron Automata Processor," in *Proceedings of the ACM International Conference on Computing Frontiers*, ser. CF '16. New York, NY, USA: ACM, 2016, pp. 135–144. [Online]. Available: http://doi.acm.org/10.1145/2903150.2903172

[14] T. Tracy, Y. Fu, I. Roy, E. Jonas, and P. Glendenning, "Towards machine learning on the Automata Processor," in *Proceedings of ISC High Performance Computing*, 2016, pp. 200–218.

[15] K. Angstadt, J. Wadden, V. Dang, T. Xie, D. Kramp, W. Weimer, M. R. Stan, and K. Skadron, "MNCaRT: An open-source, multi-architecture automata-processing research and execution ecosystem," *IEEE Computer Architecture Letters*, 2017.

[16] PCRE, "Perl compatible regular expressions," http://www.pcre.org, 2017.

[17] J. Wadden and K. Skadron, "VASim: An open virtual automata simulator for automata processing application and architecture research," University of Virginia, Tech. Rep. CS2016-03, 2016.

[18] J. Wadden, V. Dang, N. Brunelle, T. T. II, D. Guo, E. Sadredini, K. Wang, C. Bo, G. Robins, M. Stan, and K. Skadron, "ANMLZoo: a benchmark suite for exploring bottlenecks in automata processing engines and architectures," in *2016 IEEE International Symposium on Workload Characterization (IISWC)*, Sept 2016, pp. 1–12.

[19] T. Xie, V. Dang, J. Wadden, K. Skadron, and M. R. Stan, "REAPR: Reconfigurable engine for automata processing," in *Proceedings of the International Conference on Field-Programmable Logic and Applications*, 2017.

[20] J. Wadden, S. Khan, and K. Skadron, "Automata-to-Routing: An open source toolchain for design-space exploration of spatial automata processing architectures," in *Proceedings of the IEEE International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, 2017.

[21] M. Sipser, *Introduction to the Theory of Computation*. Thomson Course Technology, 2006, vol. 2.

[22] M. Becchi, "Regular expression processor," http://regex.wustl.edu, 2011.

[23] R. Cox, "Re2: a principled approach to regular expression matching," https://opensource.googleblog.com/2010/03/re2-principled-approach-to-regular.html, March 2010.

[24] X. Yu and M. Becchi, "GPU acceleration of regular expression matching for large datasets: Exploring the implementation space," in *Proceedings of the ACM International Conference on Computing Frontiers*, ser. CF '13. New York, NY, USA: ACM, 2013, pp. 18:1–18:10. [Online]. Available: http://doi.acm.org/10.1145/2482767.2482791

[25] R. Sidhu and V. K. Prasanna, "Fast Regular Expression Matching Using FPGAs," in *Proceedings of the the 9th Annual IEEE Symposium on Field-Programmable Custom Computing Machines (FCCM)*. Washington, DC, USA: IEEE Computer Society, 2001, pp. 227–238. [Online]. Available: http://dx.doi.org/10.1109/FCCM.2001.22

[26] Y. H. Yang and V. Prasanna, "High-performance and compact architecture for regular expression matching on FPGA," *IEEE Transactions on Computers*, vol. 61, no. 7, pp. 1013–1025, July 2012.

[27] E. Spishak, W. Dietl, and M. D. Ernst, "A type system for regular expressions," in *Proceedings of the 14th Workshop on Formal Techniques for Java-like Programs*, ser. FTfJP '12, 2012, pp. 20–26. [Online]. Available: http://doi.acm.org/10.1145/2318202.2318207

[28] P. Caron and D. Ziadi, "Characterization of Glushkov automata," *Theoretical Computer Science*, vol. 233, no. 1, pp. 75–90, 2000.

[29] M. Fowler, *UML Distilled: A Brief Guide to the Standard Object Modeling Language*, ser. Object Technology Series. Addison-Wesley, 2004.

[30] K. Angstadt, J. Wadden, W. Weimer, and K. Skadron, "MNRL and MNCaRT: An open-source, multi-architecture state machine research and execution ecosystem," University of Virginia, Tech. Rep. CS2017-01, 2017.

[31] Internet Engineering Task Force, *JSON Schema: core definitions and terminology*, Jan. 2013, no. json-schema-core. [Online]. Available: http://json-schema.org/latest/json-schema-core.html

[32] M. Becchi and P. Crowley, "Extending finite automata to efficiently match perl-compatible regular expressions," in *Proceedings of the ACM International Conference on emerging Networking EXperiments and Technologies*, ser. CoNEXT '08, 2008, pp. 25:1–25:12.

[33] K. Angstadt, W. Weimer, and K. Skadron, "RAPID programming of pattern-recognition processors," in *Proceedings of the 21st International Conference on Architectural Support for Programming Languages and Operating Systems*, ser. ASPLOS '16, 2016, pp. 593–605.

[34] D. Kramp, J. Wadden, and K. Skadron, "Automata Lab: An open-source automata visualization, simulation, and manipulation tool," University of Virginia, Tech. Rep. CS2017-03, 2017.

[35] V. Betz and J. Rose, "VPR: A new packing, placement and routing tool for FPGA research," in *Proceedings of the International Workshop on Field Programmable Logic and Applications*. Springer, 1997, pp. 213–222.