

OpenNWA: A Nested-Word-Automaton Library

Evan Driscoll¹, Aditya Thakur¹, and Thomas Reps^{1,2}

¹Computer Sciences Department, University of Wisconsin – Madison
{driscoll,adi,reps}@wisc.edu

²GammaTech, Inc., Ithaca, NY

Abstract. Nested-word automata (NWAs) are a language formalism that helps bridge the gap between finite-state automata and push-down automata. NWAs can express some context-free properties, such as parenthesis matching, yet retain all the desirable closure characteristics of finite-state automata.

This paper describes OpenNWA, a C++ library for working with NWAs. The library provides the expected automata-theoretic operations, such as intersection, determinization, and complementation. It is packaged with WALi—the *Weighted Automaton Library*—and interoperates closely with the weighted pushdown system portions of WALi.

1 Introduction

Many problems in computer science are solved by modeling a component as an automaton. Traditionally, this means either confronting several undecidable problems that arise with the use of pushdown automata or giving up expressivity, and usually precision, by using finite-state automata.

Recently, the development of nested word automata (NWAs) and related formalisms [2, 3] has revealed a fertile middle ground between these two extremes. NWAs are powerful enough to express some “context-free”-style properties, such as parenthesis matching, and yet retain the decidability properties that make it convenient to work with regular languages. In particular, NWAs and their languages are closed under operations such as intersection and complementation.

NWAs have been applied in areas such as modeling programs and XML documents. When modeling programs, NWAs can ensure that calls and returns match, thus eliminating spurious data flows along invalid paths. In XML documents, NWAs can model the matching between opening and closing tags.

We have created a C++ implementation of NWAs called OpenNWA. OpenNWA is packaged with the *Weighted Automata Library*, WALi [7]. WALi also provides implementations for weighted finite-state automata and weighted pushdown systems (WPDSs). The OpenNWA library

- implements (in the terminology of [3]) linearly-accepting, weakly-hierarchical nested-word automata that fully support pending calls and returns.
- provides the standard automata-theoretic operations, such as intersection and complement, except for minimization. (See §3.1 for a list and discussion.)
- is extensible via a mechanism that allows the user to attach arbitrary client information to each node in the automaton. (See §3.2.)
- inter-operates with WALi’s WPDS library and allows the user to issue queries about an NWA’s configuration space. (See §3.3.)

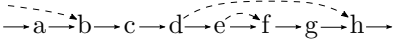
- provides extensive documentation [5] and a test suite.
- provides an NWA file format, readers, writers, and a set of command-line utilities for operating on file representations [5, §5].

OpenNWA is currently used by three projects: McVeto [11], PCCA [4], and JAM [6]. We will discuss these uses in more detail in section §4. OpenNWA can be downloaded from <http://research.cs.wisc.edu/wpis/OpenNWA>.

2 NWAs

This section describes nested-word automata [2] and related terms at an intuitive level, and gives an example of how they are used in program analysis. For the formal definitions that we use, see our technical report [5, App. A].

A nested word is an ordinary (linear) string of symbols over some alphabet Σ paired with a *nesting relation*. The nesting relation describes a hierarchical relation between input positions, for instance between matched parentheses.

Graphically, a nested word can be depicted as illustrated to the right. In this image,  following just the horizontal arrows illustrates the linear word, while the curved edges (“*nesting edges*”) indicate positions that are related by the nesting relation. For a nesting relation to be valid, nesting edges must only point forward in the word and may not share a position or cross.

Positions in the word that appear at the left end of a nesting edge are called *call positions*, those that appear at the right end are called *return positions*, and the remaining are *internal positions*. It is possible to have *pending* calls and returns, which are not matched within the word itself. For a given return, the source of the incoming nesting edge is called that return’s *call predecessor*.

Nested-word automata (NWAs) are a generalization of ordinary finite-state automata (FA). An NWA’s transitions are split into three sets—call transitions, internal transitions, and return transitions. Call and internal transitions are defined the same as transitions in ordinary FAs, while return transitions also have a call-predecessor state as an additional element.

To understand how an NWA works, consider first the case of an ordinary FA M . We can think of M ’s operation as labeling each edge in the input word with the state that M is in after reading the symbol $\xrightarrow{q_0} a \xrightarrow{q_1} b \xrightarrow{q_2} c \rightarrow d \rightarrow e$ at that edge’s source. For instance, shown to the right $\xrightarrow{q_0} a \xrightarrow{q_1} b \xrightarrow{q_2} c \rightarrow d \rightarrow e$ is a partial run. To find the next state, M looks for a transition out of q_2 with the symbol c —say with a target of q_3 —and labels the next edge with q_3 .

The operation of an NWA proceeds in a fashion similar to a standard FA, except that the machine also labels the nesting edges. When the NWA reads an internal position, it chooses a transition and labels the next linear edge the same way an FA would. When the NWA reads a call position, it picks a matching call transition and labels the next linear edge in the same way, but also labels the outgoing *nesting* edge with the state that the NWA is leaving. When the NWA reads a return position, it looks not only at the preceding linear state but also at the state on the incoming nesting edge. It then chooses a return transition that matches both, and labels the next linear edge with the target state. An example NWA run is shown in Fig. 1.




Fig. 1. An example program, corresponding NWA, and accepted word. State labels are arbitrary; transition symbols give the line number of the corresponding statement. Some nodes are elided. Dotted lines indicate call transitions, and dashed lines are return transitions. The state in parentheses on a return transition is the call predecessor.

OpenNWA supports ε internal transitions, which operate in an analogous way to ε rules in ordinary FAs. It also supports something we call *wild* transitions, which match any single symbol. Wilds can appear on any transition type.

2.1 Example

NWAs can be used to encode the interprocedural control-flow graph (ICFG) of a program. Intraprocedural ICFG edges become internal transitions, interprocedural call edges become call transitions, and interprocedural return edges become return transitions. For an ICFG return edge (*exit-site*, *return-site*), we use the call site that corresponds to *return-site* in the call-predecessor position of the NWA’s transition. The symbols on a transition depend on the application, but frequently are the corresponding statement.

An example program, the corresponding NWA, and an example word accepted by that NWA are shown in Fig. 1. Using an NWA allows us to exclude paths such as $1-2-5-6-7-8-9-\dots$ that do not follow a matched path of calls and returns. Fewer paths can allow a client analysis to obtain increased precision.

3 The OpenNWA library

OpenNWA provides a C++ class called `NWA` for representing NWAs. For information about constructing NWAs and actual API information, see the OpenNWA documentation [5]. In this section, we briefly describe some things a user can do with an NWA (or NWAs) once it is built.

3.1 Automata-theoretic operations

As mentioned in the introduction, OpenNWA supports most automata-theoretic operations:

- intersection
- union
- Kleene star
- reversal
- concatenation
- determination
- complement
- emptiness checking
- example word generation

For the most part, we use Alur and Madhusudan’s algorithms [2, 3]; however we note three exceptions. First, we implemented emptiness checking and example generation using WPDSs, discussed momentarily. Second, we found and

corrected a minor error in Kleene star [5, §6.4]. Third, we expressed Alur and Madhusudan’s determinization algorithm using relational operators [5, App. B].

OpenNWA supports determining whether an NWA’s language is empty, and if it is not, OpenNWA can return an arbitrary word from the NWA’s language. It is also possible to ask specifically for a shortest accepted word.

The library performs these operations by converting the NWA into a WPDS and running a poststar query from the set of initial configurations (see §3.3). To find an example word, OpenNWA uses a witness-tracing version of poststar [10, §3.2.1], and extracts the word from the resulting witness. For the shortest word, we simply use weights that track the length of the path from the initial state.

The ability to obtain an example word is important in program analysis. It can show the end user of an analysis tool a program trace that may violate a property. Moreover, this feature is fundamental to counterexample-guided refinement: in CEGAR-based model checkers, the counterexample is typically an example word from the automaton that models the program.

3.2 Client information

OpenNWA provides a facility that we call *client information*. This feature allows the user of the library to attach arbitrary information to each state in the NWA. For instance, as discussed in §4, McVeto uses NWAs internally, and uses client information to attach a formula to each state in the NWA.

The library tracks this information as best as it can through each of the operations discussed in the previous section, and supports callback functions to compute new client information when it does not have the information it needs.

3.3 Inter-operability with WPDSs

Weighted pushdown systems (WPDSs) can be used to perform interprocedural static analysis of programs [10]. The PDS proper provides a model of the program’s control flow, while the weights on PDS rules express the dataflow transformers. Algorithms exist to query the configuration space of WPDSs, which correspond to asking a question about the data values that can arise at a set of configurations in the program’s state space. A configuration consists of a control location and a list of items on the stack.

OpenNWA supports converting an NWA into a WPDS implemented by WALi. This feature allows an OpenNWA client to issue queries about the configuration space of an NWA. (For instance, our `isLanguageEmpty()` function effectively asks a query of the form “Is it possible to start in an initial configuration and consume a nested word to reach a configuration where the automaton is in an accepting state?”) The WPDS’s stack corresponds to the states that label the nesting edges in an input nested word.

NWAs themselves are not weighted, but the library provides a facility for determining the weights of the WPDS rules during the conversion. The user provides an instance of a subclass of `WeightGen`, which acts as a factory function. The function is called with the states in question and return the weight of the resulting WPDS rule. It is, of course, possible to use the client information of the states in question to determine the weight.

4 Uses of the library

In this section, we briefly discuss the current users of OpenNWA.

I/O compatibility checking. We used OpenNWA as the primary component of a tool called the Producer-Consumer Conformance Analyzer (PCCA) [4]. Given two programs that operate in a producer/consumer relationship over a stream, PCCA’s goal is to determine whether the consumer is prepared to accept all messages that the producer can emit, and find a counterexample if not. It proceeds by inferring a model of the output language of the producer, inferring a model of the input language of the consumer, and determining whether the models are compatible.

We used NWAs for our models, building them from the ICFG as discussed in §2.1. Edges corresponding to statements that perform I/O are labeled with the type of the I/O, and all other internal transitions are labeled with ε . Conceptually what we want to check is whether the producer’s output language is a subset of the consumer’s input language, which is an operation NWAs and our library support. In reality this check is likely to be too strict, and we need an additional step, detailed in [4, §2.3 and §3.2].

Machine-code model checking. McVeto is a machine-code verification engine [11] that, given a binary and description of a bad target state, tries to find either (i) an input that forces the bad state to be reached, or (ii) a proof that the bad state is impossible to reach.

McVeto uses an model of the program called a *proof graph*, which is an NWA that overapproximates the program’s behavior. States in a proof graph are labeled with formulas; edges are labeled with program statements. McVeto starts with a very coarse overapproximation, which it then refines. One refinement technique uses symbolic execution to produce a concrete trace of the program’s behavior, performs *trace generalization* [11, §3.1] to convert the trace into an overapproximating NWA (the “folded trace”), and intersects the current proof graph with the folded trace to obtain the new proof graph. The formula on a state in the new proof graph is the conjunction of the formulas on the states that are being paired from the current proof graph and the folded trace.

McVeto’s implementation uses OpenNWA’s client-information feature to store the formula for each state. During intersection, the callback functions mentioned in §3.2 compute the conjunction of the input formulas, which are then used in the new proof graph.

To determine whether the target state is not reachable in the proof graph (and thus is definitely not reachable in the actual program), McVeto calls `prestar()` (see §3.3). The result of this call is also used to determine which “frontier” to extend next during directed test generation [11, §3.3].

JavaScript security-policy checking and weaving. The JAM tool [6] checks a JavaScript program against a security policy, either verifying that the program is already correct with respect to that policy or inserting dynamic checks into the program to ensure that it will behave correctly. JAM builds *two* models of the input program, one that overapproximates the control flow of the

program and one that overapproximates the data flow. The policy is also expressed as an NWA of forbidden traces. By intersecting the policy automaton with both program models, JAM obtains an NWA that expresses traces that possibly violate the policy.

Once JAM has the combined NWA, it asks OpenNWA for a shortest word in the language. If the language is empty (i.e., there is no shortest word), the program always respects the policy. If OpenNWA returns an example word w , JAM checks whether w corresponds to a valid trace through the program. If w is valid, then JAM inserts a dynamic check to halt concrete executions corresponding to w that would violate the policy. If w is not valid, then JAM can either refine the abstraction and repeat, or insert a dynamic check to detect and halt concrete executions corresponding to w for which the policy would be violated.

5 Related work

Alur and Madhusudan each maintain a page giving a significant bibliography of papers that present theoretical results, practical applications, and tools related to NWAs and visibly pushdown automata (VPAs) [1, 8]. VPAs and their languages are another formalism which can be seen as an alternative encoding of NWAs and nested-word languages [3].

VPALib [9] is a general-purpose library implementing VPAs. However, OpenNWA's implementation is far more complete. For instance, VPALib does not support concatenation, complementation (although it does support determinization), checking emptiness, or obtaining an example word.

References

1. R. Alur. Nested words, 2011. <http://www.cis.upenn.edu/~alur/nw.html>.
2. R. Alur and P. Madhusudan. Adding nesting structure to words. In *DLT*, 2006.
3. R. Alur and P. Madhusudan. Adding nesting structure to words. *JACM*, 56(3), May 2009.
4. E. Driscoll, A. Burton, and T. Reps. Checking conformance of a producer and a consumer. In *FSE*, 2011.
5. E. Driscoll, A. Thakur, A. Burton, , and T. Reps. WALi: Nested-word automata. TR-1675R, Comp. Sci. Dept., Univ. of Wisconsin, Madison, WI, Sept. 2011.
6. M. Fredrikson, R. Joiner, S. Jha, T. Reps, P. Porras, H. Saidi, and V. Yegneswaran. Efficient runtime policy enforcement using counterexample-guided abstraction refinement. Under submission to CAV 2012.
7. N. Kidd, A. Lal, and T. Reps. WALi: The Weighted Automaton Library, 2007. www.cs.wisc.edu/wpis/wpds/download.php.
8. P. Madhusudan. Visibly pushdown automata – automata on nested words, 2009. <http://www.cs.uiuc.edu/~madhu/vpa/>.
9. H. Nguyen. Visibly pushdown automata library, 2006. <http://www.emn.fr/z-info/hnguyen/vpa/>.
10. T. Reps, S. Schwoon, S. Jha, and D. Melski. Weighted pushdown systems and their application to interprocedural dataflow analysis. *SCP*, 58(1–2), Oct. 2005.
11. A. Thakur, J. Lim, A. Lal, A. Burton, E. Driscoll, M. Elder, T. Andersen, and T. Reps. Directed proof generation for machine code. TR 1669, UW-Madison, Apr. 2010. Abridged version published in CAV 2010.

A Reviewers' Appendix

In this appendix we give some technical material that is omitted from the paper proper, but present in our technical report [5].

A.1 Definitions

The definitions below match those of Alur and Madhusudan's original definition [2], which is what our implementation is based off of. Later work revised the definition of NWAs so that the automata are more complex [3]; the revised definitions recognize the same class of languages, but may be more concise than what we support. (However, most of the closure operations cannot operate on the expanded version directly.)

We start off with a formal definition of nested words and nested-word languages.

Definition 1. A *nested word* (w, \rightsquigarrow) over alphabet Σ is an ordinary (linear) word $w \in \Sigma^*$ of length $|w|$ together with a *nesting relation* \rightsquigarrow . The nesting relation is a collection of edges (over the positions in w) that do not cross. Formally, $\rightsquigarrow \subseteq \{-\infty, 1, 2, \dots, |w|\} \times \{1, 2, \dots, |w|, +\infty\}$ such that:

- Nesting edges only go forward: if $i \rightsquigarrow j$ then $i < j$.
- No two edges share a position: for $1 \leq i \leq |w|$, there is at most one j such that $i \rightsquigarrow j$ or $j \rightsquigarrow i$.
- Edges do not cross: if $i \rightsquigarrow j$ and $i' \rightsquigarrow j'$, then one cannot have $i < i' \leq j < j'$.

When $i \rightsquigarrow j$ holds, for $1 \leq i \leq |w|$, i is called a *call* position. If $i \rightsquigarrow +\infty$, then i is a *pending call*; otherwise i is a *matched call*, and the (unique) position j such that $i \rightsquigarrow j$ is called its *return successor*. (Note that these terms refer to positions within w rather than the symbol, and one symbol from Σ can be used in more than one kind of position.)

Similarly, when $i \rightsquigarrow j$ holds, for $1 \leq j \leq |w|$, j is a *return* position. If $-\infty \rightsquigarrow j$, then j is a *pending return*, otherwise j is a *matched return*, and the (unique) position i such that $i \rightsquigarrow j$ is called its *call predecessor*.

A position $1 \leq i \leq |w|$ that is neither a call nor a return is an *internal* position.

Definition 2. A *nested-word language* is a set of nested words; such a language is a *regular nested-word language* if it is accepted by an NWA as defined below. However, the term *nested-word language* is commonly abused to refer to only regular nested-word languages; we continue this tradition.

Now we can formally define NWAs. We will start by excluding transitions with ε or wild, and then expand upon this definition.

Definition 3. A *nested-word automaton* (NWA) A is a 5-tuple $(Q, \Sigma, Q_0, \delta, F)$, where Q is a finite set of states, Σ is a finite alphabet, $Q_0 \subseteq Q$ is a set of initial states, $F \subseteq Q$ is a set of final states, and δ is a transition relation. The transition relation δ consists of three components, $(\delta_c, \delta_i, \delta_r)$, where:

- $\delta_i : (Q \times \Sigma) \times Q$ gives *internal transitions*
- $\delta_c : (Q \times \Sigma) \times Q$ gives the *call transitions*
- $\delta_r : (Q \times Q \times \Sigma) \times Q$ gives the *return transitions*

Starting from a state $q_0 \in Q_0$, an NWA A reads a nested word (w, \rightsquigarrow) from left to right, and performs transitions according to the current input symbol and \rightsquigarrow . If A is in state q when reading input symbol σ at position i in w , and i is an internal (resp, call)

position, A makes a transition to a state q' (if one is available) such that $(q, \sigma, q') \in \delta_i$ (resp, $(q, \sigma, q') \in \delta_c$). If i is a return position, let k be the call predecessor of i (so $k \rightsquigarrow i$) and q_c be the state that A was in just before the transition it made on the k^{th} symbol; A changes to a state q' such that $(q, q_c, \sigma, q') \in \delta_r$. If there is a computation of A on input (w, \rightsquigarrow) that terminates in a state $q \in F$, then A accepts (w, \rightsquigarrow) .

Operationally it is sometimes beneficial to think of an NWA by analogy to a restricted pushdown automaton. The PDA's stack alphabet is the same as the set of states, and the behavior is as follows:

- When reading a call, the PDA pushes (just) the current state onto the stack
- When reading an internal position, the PDA may not modify the stack
- When reading a return, the PDA must pop exactly one item from the stack.

This definition is roughly that of a related model called a visibly pushdown automaton (VPA) [3], and VPAs and NWAs and their languages can be seen as alternative encodings of each other.

We extend the definition of NWAs give above to support ε and wild transitions, for which we use the symbol $@$. Internal transitions are the only type allowed to have ε , and the definition is extended as expected:

- $\delta_i \subseteq (Q \times (\Sigma \cup \{\varepsilon\})) \times Q$
- While an NWA A is reading a nested word, at any point A is allowed to select a transition $(q, \varepsilon, q') \in \delta_i$, where q is A 's current state, and change to state q' .


Wild transitions match any single input character, and can appear on any kind of transition. Thus if we let $\Sigma_{@}$ be $\Sigma \cup \{@\}$, we have

- $\delta_i \subseteq (Q \times (\Sigma_{@} \cup \{\varepsilon\})) \times Q$
- $\delta_c \subseteq (Q \times \Sigma_{@}) \times Q$
- $\delta_r \subseteq (Q \times Q \times \Sigma_{@}) \times Q$

The operation of the NWA is modified as follows. When reading a nested word and consuming the input σ , an NWA A is allowed to select a transition of the form $(q, @, q')$ or $(q, q_c, @, q')$ instead of (q, σ, q') or (q, q_c, σ, q') ; the type of transition (call, internal, return) has to match the current input position.

A.2 Kleene-star mistake

As mentioned in §3, we found a minor error in Alur and Madhusudan's formulation of Kleene star [3, theorem 3.6]. In effect, the error is analogous to what happens in Thompson's construction for standard finite automata if one does not add a new distinguished start state, but instead just connects the accepting states to the initial states with ε transitions and makes the initial states accepting (to accept ε), as illustrated in the diagram to the right. If there is a cycle that allows the automaton to return to the initial state, that path will give a string that should not be accepted. Alur confirmed our diagnosis.



Adding a distinguished start state, as detailed in our technical report [5, §6.4], fixes the problem.

A.3 NWA-to-WPDS conversion

§3.3 discusses our conversion from NWA to WPDSs. In this section, we give details. We first define pushdown systems (PDSs), then define our conversion, and finally use the example from Fig. 1 to illustrate the conversion.


Definition 4. A **pushdown system** (PDS) is a three-tuple $\mathcal{P} = (P, \Gamma, \Delta)$, where P is a finite set of **control locations**, Γ is a finite set of **stack symbols**, and $\Delta \subseteq P \times \Gamma \times P \times \Gamma^*$ is a finite set of **rules**. A **configuration** of \mathcal{P} is a pair $\langle p, u \rangle$ where $p \in P$ and $u \in \Gamma^*$. A rule $r \in \Delta$ is written as $\langle p, \gamma \rangle \hookrightarrow \langle p', u \rangle$, where $p, p' \in P$, $\gamma \in \Gamma$, and $u \in \Gamma^*$. The rules define a **transition relation** \Rightarrow on configurations of \mathcal{P} as follows: If $r = \langle p, \gamma \rangle \hookrightarrow \langle p', u' \rangle$, then $\langle p, \gamma u \rangle \Rightarrow \langle p', u' u \rangle$ for all $u \in \Gamma^*$.

Definition 5. Given an NWA $A = (Q, \Sigma, Q_0, \delta, F)$, we define PDS $\mathcal{P}_A = (P, Q, \Delta)$. The set of control locations P is defined as $\{s\} \cup \{s_q \mid q \in Q\}$. Each transition of A is converted to one or two rules in Δ , as follows:

- For each transition $((q, \sigma), q') \in \delta_i$, Δ has a rule $\langle s, q \rangle \hookrightarrow \langle s, q' \rangle$.
- For each transition $((q, \sigma), q') \in \delta_c$, Δ has a rule $\langle s, q \rangle \hookrightarrow \langle s, q' q \rangle$. (The PDS's transition system will leave q on the stack and push q' .)
- For each transition $((q, q_c, \sigma), q') \in \delta_r$, Δ has two rules, $\langle s, q \rangle \hookrightarrow \langle s_q, \varepsilon \rangle$ and $\langle s_q, q_c \rangle \hookrightarrow \langle s, q' \rangle$. (Conceptually this can be thought of as a single transition $\langle s, q q_c \rangle \hookrightarrow \langle s, q' \rangle$ of a prefix rewriting system.)

One can interpret this conversion as simply moving the information in the NWA's state into the top symbol of the stack: then, e.g., internal moves can change the top symbol (but do no more).

We now will illustrate the conversion on our example, reproduced below from Fig. 1. We list all the tuples in the transition relation δ , and give the conversion to PDS rules.



	NWA transition	PDS rule(s)
δ_i	$((A, 1), B)$	$\langle s, A \rangle \hookrightarrow \langle s, B \rangle$
	$((C, 3), D)$	$\langle s, C \rangle \hookrightarrow \langle s, D \rangle$
	$((E, 6), F)$	$\langle s, E \rangle \hookrightarrow \langle s, F \rangle$
	$((E, 6), G)$	$\langle s, E \rangle \hookrightarrow \langle s, G \rangle$
	$((F, 7), K)$	$\langle s, F \rangle \hookrightarrow \langle s, K \rangle$
	$((H, 9), K)$	$\langle s, H \rangle \hookrightarrow \langle s, K \rangle$
δ_c	$((B, 2), E)$	$\langle s, B \rangle \hookrightarrow \langle s, EB \rangle$
	$((G, 8), E)$	$\langle s, G \rangle \hookrightarrow \langle s, EG \rangle$
δ_r	$((K, B, 10), C)$	$\langle s, K \rangle \hookrightarrow \langle s_K, \varepsilon \rangle$ $\langle s_K, B \rangle \hookrightarrow \langle s, C \rangle$
	$((K, G, 10), H)$	$\langle s, K \rangle \hookrightarrow \langle s_K, \varepsilon \rangle$ $\langle s_K, G \rangle \hookrightarrow \langle s, H \rangle$