

*Clingraph: A System for ASP-based Visualization**

SUSANA HAHN

University of Potsdam, Potsdam, Germany Potassco Solutions, Potsdam, Germany
(e-mail: susana.hahn.martin.lunas@uni-potsdam.de)

ORKUNT SABUNCU

TED University, Ankara, Turkey Potassco Solutions, Potsdam, Germany
(e-mail: orkunt.sabuncu@tedu.edu.tr)

TORSTEN SCHAUB and TOBIAS STOLZMANN

University of Potsdam, Potsdam, Germany Potassco Solutions, Potsdam, Germany
(e-mails: torsten.schaub@uni-potsdam.de, tobias.stolzmann@uni-potsdam.de)

submitted 22 December 2022; revised 24 November 2023; accepted 08 January 2024

Abstract

We present the Answer Set Programming (ASP)-based visualization tool *clingraph*, which aims at visualizing various concepts of ASP by means of ASP itself. This idea traces back to the *aspviz* tool and *clingraph* redevelops and extends it in the context of modern ASP systems. More precisely, *clingraph* takes graph specifications in terms of ASP facts and hands them over to the graph visualization system *graphviz*. The use of ASP provides a great interface between logic programs and/or answer sets and their visualization. Also, *clingraph* offers a Python application programming interface (API) that extends this ease of interfacing to *clingo*'s API and in turn to connect and monitor various aspects of the solving process.

KEYWORDS: logic programming methodology and applications, knowledge representation and nonmonotonic reasoning

1 Introduction

With the advance of Answer Set Programming (ASP; Lifschitz 2019) into more and more complex application domains, also the need for inspecting problems as well as their solution increases significantly. The intrinsic difficulty lies in the fact that ASP constitutes a general problem solving paradigm, whereas the wide spectrum of applications rather calls for customized presentations.

We address this by taking up the basic idea of *aspviz* (Cliffe *et al.* 2008), to visualize ASP by means of ASP itself, and extend it in the context of modern ASP systems. The resulting system is called *clingraph* (v1.1.0).^{1,2} The common idea is to specify a

* This paper is an extended version of an article presented at LPNMR'22 ((Hahn *et al.* 2022)).

¹ <https://github.com/potassco/clingraph>.

² <https://clingraph.readthedocs.io>.

```

2  node(1..6).
3  edge(1,2). edge(1,3). edge(1,4). edge(2,4). edge(2,5).
4  edge(2,6). edge(3,4). edge(3,5). edge(5,6).
5  color(red; green; blue).

6  { assign(N, C) : color(C) } = 1 :- node(N).
7  :- edge(N, M), assign(N, C), assign(M, C).

9  #show node/1.
10 #show edge((N,M)) : edge(N, M).
11 #show attr(graph_nodes, default, style, filled).
12 #show attr(node, N, color, C) : assign(N, C).

```

Listing 1. *Graph coloring instance, encoding, and display (color.lp).*

visualization in terms of a logic program that defines special atoms capturing graphic elements. This allows us to customize the presentation of an application domain by means of ASP and thus to easily connect with the problem specification and its solutions.

The visualization in *clingraph* rests upon graph structures that are passed on to the graph layout system *graphviz*.³ *Graphviz* layout programs process graph descriptions written in a simple text language. We employ their python application programming interface (API) to generate this description from an ASP representation. To this end, *clingraph* takes – in its basic setting – a set of facts over predicates **graph**, **node**, **edge**, and **attr** as input and produces an output visualizing the induced graph structure.

As a simple example, consider the graph coloring problem in Listing 1.⁴ The actual problem instance and encoding are given in Lines 1–4 and 6–7, respectively. However, of particular interest are Lines 9–12 that use **#show** directives to translate the resulting graph colorings into *clingraph*'s input format. While Lines 9 and 10 account for the underlying graph, the two remaining lines comprise instructions to *graphviz*. Line 11 fixes the layout of graph nodes. More interestingly, Line 12 translates the obtained graph coloring to layout instructions for *graphviz*. Here, all entities are grouped under a default graph labeled **default** (which can be changed via an option; similarly, graphs are taken to be undirected unless changed by option **--type**).

Launching *clingo* so that only the resulting stable model is obtained as a set of facts allows us to visualize the result via *clingraph*:

```

clingo --outf=0 -V0 --out-atomf=%s. color.lp | head -n1 | \
clingraph --out=render --format=png

```

The used options suppress *clingo* output and transform atoms into facts; the intermediate UNIX command extracts the line comprising the stable model. Note that one can also use a solver other than *clingo* to generate the stable model in the expected form. The final call to *clingraph* produces a file in PNG format, shown in Figure 1.

Obviously, the above proceeding only reflects the very basic functionality of *clingraph*. We elaborate upon its extended functionality in the next section and present a series of illustrative cases studies in Section 3. They range from the visualization of stable models, including animated dynamic solutions as well as interactive ones, over visualizing the actual solving process, to the visual inspection of the structure of logic programs. Finally,

³ <https://graphviz.org>.

⁴ In this and other Listings, we include the name of the respective file in parentheses. This information serves as a reference point when presenting command-line instructions.

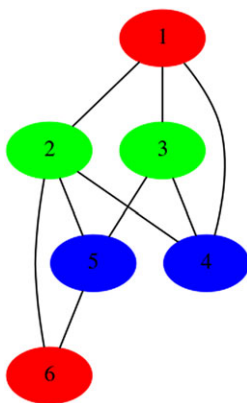


Fig. 1. Visualization of the (first) stable model of the logic program in Listing 1.

we present in Section 4 some support for generating strings by using a template engine. We relate our approach to others' and summarize it in Sections 5 and 6, respectively.

2 *Clingraph*

In its most basic setting, *clingraph* can be regarded as a front-end to *graphviz* that relies on the fact format sketched above. In fact, the full-fledged version of the fact format allows for specifying multiple graphs as well as subgraphs. The former is done by supplying several instances of predicate `graph/1` whose only argument provides an identifier for regrouping all elements belonging to the graph at hand. To that effect, there are also binary versions of predicates `node` and `edge`, whose second argument refers to the encompassing graph. For example, the following facts describe `n` graphs, each with one edge connecting two nodes.

```

1 id(1..n).
2 graph(g(X)) :- id(X).
3 node(n((a;b), X),g(X)) :- id(X).
4 edge((n(a,X),n(b,X)),g(X)) :- id(X).

```

Multiple graphs are of particular interest when visualizing dynamic domains, as in planning, where each graph may represent a state of the world. We illustrate this in Section 3 and show how the solution to a planning problem can be turned into an animation.

Subgraphs⁵ are specified by the binary version of predicate `graph`, whose second argument indicates the super-ordinate graph. For instance, replacing Line 2 above by the following two rules makes `g(X)` a subgraph of `g(X+1)` for `X=1..n-1`.

```

graph(g(X))          :- id(X), not id(X+1).
graph(g(X),g(X+1)) :- id(X),   id(X+1).

```

Clingraph allows for selecting designated graphs by supplying their identifier to option `--select-graph`; several ones are selected by repeating the option with the respective identifiers on the command line.

As exemplified in Listing 1, the quaternary predicate `attr/4` describes properties of graph elements; this includes all attributes of *graphviz*. The first argument fixes the type

⁵ Subgraphs correspond to clusters in *graphviz*.

of the element, namely, `graph`, `node`, and `edge`, along with keywords `graph_nodes` and `graph_edges` to refer to all nodes and edges of a graph, respectively. The second argument gives the identifier of the element, and the last two provide the name and value of the *graphviz* attribute. Some attributes, mainly labels, are often constructed by concatenating multiple values. We simplify this burden by providing an integration with a template engine to allow string formatting. In Section 4, we describe this extension in detail.

In order to avoid name clashes, *clingraph* offers the option `--prefix` to change all graph-oriented predicates by prepending a common prefix. For instance, `--prefix='viz-'` changes the dedicated predicate names to `viz-graph`, `viz-node`, `viz-edge`, and `viz-attr` while maintaining their arities.

The more interesting use cases emerge by using *visualization encodings*. While in our introductory example the latter was mimicked by `#show` statements, in general, a visualization encoding can be an arbitrary logic program producing atoms over the four graph-oriented predicates. Obviously, when it comes to visualization, a given problem encoding can then be supplemented with a dedicated visualization encoding, whose output is then visualized by *clingraph* as shown in the introductory section.

In practice, however, it turns out that this joint approach tends to slow down the solving process significantly. Rather, it is often advantageous to resort to a sequential approach, in which the stable models of the problem encoding are passed to a visualization encoding. This use case is supported by *clingraph* with extra functionality when using the ASP system *clingo*. More precisely, this functionality relies upon the *clingo* feature to combine the output of a run, possibly comprising various stable models, in a single `json` object.⁶ To this end, *clingraph* offers the option `--select-model` to select one or multiple stable models from the `json` object. Multiple models are selected by repeating the option with the respective number.

To illustrate this, let us replace Line 1 above by

```
{ id(1..n) } = 1.
```

to produce `n` stable models with one graph each, rather than a single model with `n` graphs as above. The handover of all stable models of the resulting logic program in `multiple.lp` to *clingraph* can then be done by the following command:

```
clingo --outf=2 -c n=10 0 multiple.lp | \
clingraph --out=tex --select-model=0 --select-model=9
```

The option `--outf=2` instructs *clingo* to produce a single `json` object as output. We request all 10 stable models via “`-c n=10 0.`” Then, *clingraph* produces a \LaTeX file depicting the graphs described in the first and tenth stable model.

In the quite frequent case that the stable models are produced exclusively by the problem encoding, an explicit visualization encoding can be supplied via option `--viz-encoding` to make *clingraph* internally produce the graphic representation from the given stable models employing the *clingo* API. To ease the development of visualization encodings, *clingraph* also provides a set of external Python functions (see Section 3 for an example).

Just like *clingraph*'s input, also its output may consist of one or several graph representations. The specific representation is controlled by option `--out` that can take the following values:

⁶ <https://www.json.org>.

- **facts** produces the facts obtained after preprocessing (default)
- **dot** produces graph representations in the language DOT
- **render** generates images with the rendering method of *graphviz*
- **animate** generates a GIF after rendering
- **tex** produces a \LaTeX file

The default option **facts** allows us to inspect the processed input to *clingraph* in fact format. This involves the elimination of atoms irrelevant to *clingraph* as well as the normalization of the graph representation (e.g., turning unary predicates **node** and **edge** into binary ones, etc.). Options **dot** and **tex** result in text-based descriptions of graphs in the languages DOT and \LaTeX . These formats allow for further postprocessing and editing upon document integration. The \LaTeX file is produced with *dot2tex*.⁷ Arguments to *dot2tex* can be passed through *clingraph* via **--tex-param**. At long last, the options **render** and **animate** synthesize images for the graphs at hand. While the former aims at generating one image per graph, the latter allows us to combine several graphs in an animation. The format of a rendered graph is determined by option **--format**. The option defaults to PDF and alternative formats include PNG and SVG (cf. Section 3.3). Animation results in a GIF file. It is supported by options **--fps** to fix the number of frames per second and **--sort** to fix the order of the graphs' images in the resulting animation. The latter provides a handful of alternatives to describe the order in terms of the graph identifiers.

Also, it is worth mentioning that *clingraph*'s option **--engine** allows us to choose among the eight layout engines of *graphviz*;⁸ it defaults to **dot** which is optimized for drawing directed graphs.

Last but not least, *clingraph* also offers an API for Python. Besides *graphviz*, it heavily relies on *clorm*,⁹ a Python library providing an Object Relational Mapping (ORM) interface to *clingo*. Accordingly, the major components of *clingraph*'s API are its **Factbase** class, providing functionality for manipulating sets of facts via *clorm*, and the **graphviz** package, gathering functionality for interfacing to *graphviz*. We refer the interested reader to the API documentation for further details.¹⁰ In conjunction with *clingo*, the API can be used for visualizing the solving process. Two natural interfaces for this are provided by the **on_model** callback of *clingo*'s **solve** method as well *clingo*'s **Propagator** class. For example, the former would allow for visualizing the intermediate stable models obtained when converging to an optimal model during optimization. The latter provides an even more fine-grained approach that allows for monitoring the search process by visualizing partial assignments (cf. Section 3.4).

3 Case studies

In this section, we list several case studies as examples that showcase various features of *clingraph*. Our first example, in Section 3.1, visualizes a solution of the well-known Queens puzzle. The next one, in Section 3.2, aims at visualizing a dynamic problem of a robotic

⁷ <https://dot2tex.readthedocs.io>.

⁸ <http://www.graphviz.org/docs/layouts>.

⁹ <https://github.com/potassco/clorm>.

¹⁰ <https://clingraph.readthedocs.io/en/latest/clingraph/api.html>.

```

1 1 { queen(I,1..n) } 1 :- I = 1..n.
2 1 { queen(1..n,J) } 1 :- J = 1..n.
3   :- 2 { queen(D-J,J) }, D = 2..2*n.
4   :- 2 { queen(D+J,J) }, D = 1-n..n-1.

6 cell(1..n,1..n).

```

Listing 2. *Queens puzzle (queens.lp)*

intra-logistics scenario where the resulting animation points out temporal aspects of the problem. Then, in Section 3.3, we explore interactivity in visualizations via *clingraph*'s SVG extension. Finally, the last two case studies concentrate on visualizing aspects other than solutions of a problem. To that end, these approaches visualize the solving process of the solver (Section 3.4) and the structure of the input program (Section 3.5).

Many of these case studies need complex attributes, mainly labels, which are composed of various values. Thanks to the template engine integrated in *clingraph*, one can specify such an attribute conveniently by a template string having variables. Then, separate rules may set values of these variables in a modular way. Note that details on the usage of templates can be found in Section 4.

The interested reader is referred for further details on these examples and many others to *clingraph*'s distribution.¹¹

3.1 Visualizing a solution of the Queens puzzle

As a first example, consider the encoding of the Queens puzzle in Listing 2.¹² The idea is to place n queens on an $n \times n$ chessboard so that no two queens attack one another. A solution is captured by atoms over predicate `queen/2`. The one comprised in the first stable model of `queens.lp` for `n=5` is depicted in Figure 2. First of all, we note that the actual graph is laid out as a 5×5 grid of white and gray squares. Each atom `queen(x,y)` is then represented by putting the symbol ♔ on the square with coordinate (x,y) . All other squares are simply labeled with their actual coordinate.

The visualization encoding producing the chessboard in Figure 2 is given in Listing 3; it is used to generate the PDF in Figure 2 in the following way.

```

clingo queens.lp -c n=5 --outf=2 | \
clingraph --viz-encoding=viz.lp --out=render --engine=neato

```

To better understand the visualization encoding, it is important to realize that we use *neato* as layout engine, since it is better-suited for dealing with coordinates than the default engine *dot*.

Let us now have a closer look at the encoding in Listing 3. Interestingly, our graph consists of nodes only; no edges are provided. This is because nodes are explicitly positioned, and no edges are needed to connect them. More precisely, one node is introduced in Line 1 for each cell of the chessboard.¹³ The remainder of the encoding is concerned

¹¹ <https://github.com/potassco/clingraph/tree/master/examples>.

¹² <https://github.com/potassco/clingraph/tree/master/examples/queens>.

¹³ Strictly speaking, the definition of predicate `cell/2` belongs to the visualization encoding. Nonetheless, we add it to the problem encoding since the dimension of the board, viz. `n`, is unavailable in the visualization encoding. This is a drawback of the sequential approach: information must be shared via the stable models.






(1,5)	(2,5)		(4,5)	(5,5)
	(2,4)	(3,4)	(4,4)	(5,4)
(1,3)	(2,3)	(3,3)		(5,3)
(1,2)		(3,2)	(4,2)	(5,2)
(1,1)	(2,1)	(3,1)	(4,1)	

Fig. 2. Visualization of (first) stable model of the logic program in Listing 2.

```

1 node((X,Y)):- cell(X,Y).

3 attr(node,(X,Y),width,1):- cell(X,Y).
4 attr(node,(X,Y),shape,square):- cell(X,Y).
5 attr(node,(X,Y),style,filled):- cell(X,Y).
6 attr(node,(X,Y),fillcolor,gray):- cell(X,Y),(X+Y)\2 = 0.
7 attr(node,(X,Y),fillcolor,white):- cell(X,Y),(X+Y)\2 != 0.
8 attr(node,(X,Y),fontsize,"50"):- queen(X,Y).
9 attr(node,(X,Y),label,"♔") :- queen(X,Y).
10 attr(node,(X,Y),pos,@pos(X,Y)):- cell(X,Y).

```

Listing 3. Visualization encoding for Queens puzzle (*viz.lp*)

with the layout and positioning of each individual node, as reflected by the first and second argument of all remaining atoms over `attr/4`. This is done in a straightforward way in Lines 3–5 to fix the `width`, `shape`, and `style` of each node. Lines 6 and 7 care about the alternating coloration of nodes, depending on whether the sum of their coordinates is even or odd. The next two lines deal with cells occupied by queens. Unlike the previous rules that only refer to the problem instance, here the derived attributes depend on the obtained solution. That is, for each atom `queen(x,y)`, Line 8 fixes the `fontsize` of the `label` ♔ attributed to node (x,y) in Line 9. Whenever no `label` is given to a node, its name is used instead, as witnessed by Figure 2. Finally, Line 10 handles the positioning of nodes. In *neato*, positions are formatted by two comma-separated numbers and entered in a node’s `pos` attribute. If an exclamation mark “!” is given as a suffix, the node is also pinned down. The necessary transformation from pairs of terms is implemented by the external Python function `pos(x,y)` provided by *clingraph*. This function turns a node

identifier (x, y) into a string of form " $x, y!$ ". For each node, the result is then inserted as the fourth argument of predicate `attr/4` in Line 10.

3.2 Visualizing dynamic problems

As a second example, let us look at a dynamic problem whose solutions can be visualized in terms of animations. To this end, we have chosen a robotic intra-logistics scenario from the *asprilo* framework (Gebser et al. 2018). This scenario amounts to an extended multi-agent pathfinding problem having robots transport shelves to picking stations and back somewhere. The goal is to satisfy a batch of orders by transporting shelves covering all requested products to the picking station. For brevity, we do not reproduce the actual problem encoding¹⁴ here and rather restrict our attention to the input to the visualization encoding. The input consists of action and fluent atoms accounting for a solution and how it progresses the problem scenario over time, namely,

- `move(robot(r), (d_x, d_y), t)`¹⁵ and
- `position(o , (x, y), t)` for o among `robot(r)`, `shelf(s)`, and `station(p)`.

A `move` atom indicates that a robot r moves in the cardinal direction (d_x, d_y) at time step t (for $d_x, d_y \in \{-1, 0, 1\}$ such that $|d_x + d_y| = 1$). A `position` atom tells us that object o is at position (x, y) at time step t . All atoms sharing a common time step capture a state induced by the resulting plan.

The idea of the visualization encoding is now to depict a sequence of such states by combining the visualizations of individual states in an animation. Each state is represented by a graph that lays out the grid structure of a warehouse. We use consecutive time steps to identify and to order these graphs. This results in an atom `graph(t)` for each time step t . Similarly, we identify nodes with their coordinate along with a timestamp. This is necessary because nodes require a unique identifier across all (sub)graphs. As well, we use edges indexed by time steps to trace (the last) movements:

- `node(((x, y), t), t)`
- `edge((((x', y'), t), (($x' + d_x, y' + d_y$), t)), t)`

The first atom expresses that node $((x, y), t)$ belongs to graph t . Similarly, the second one tells us that the edge from node $((x', y'), t)$ to node $((x' + d_x, y' + d_y), t)$ belongs to graph t . It is induced by an action `move(robot(r), (d_x, d_y), t)` and its precondition `position(robot(r), (x', y'), $t - 1$)`.

Having settled the representation of graphs along with their nodes and edges, the rest of the visualization encoding mainly deals with setting their attributes.

To see this, consider Table 1, giving excerpts of the actual visualization encoding (using line numbers in the full encoding; lines in between have been dropped for brevity).¹⁶ The definition of graphs, nodes, and edges is given in Lines 19, 27, and 30–31. Let us discuss the remaining lines of interest of `viz-asprilo.lp` by inspecting some features of a visualization, produced as follows.

¹⁴ <https://github.com/potassco/asprilo-encodings>.

¹⁵ We refrain from visualizing pickup and putdown actions, and rather represent them implicitly.

¹⁶ <https://github.com/potassco/clingraph/tree/master/examples/asprilo>.

Table 1. *Selected lines from the visualization encoding for an asprilo scenario (viz-asprilo.lp).*

```

10 free(P,T) :- not position(_,P,T), position(P), step(T).
12occo(P,T,robot(R)) :- position(robot(R),P,T),
13    not position(station(_),P,T),
14    not position(shelf(_),P,T).
19 graph(T) :- step(T).
27 node((P,T),T) :- position(P), step(T).
30edge((((X,Y),T),((X+DX,Y+DY),T)),T) :- move(robot(R),(DX,DY),T),
31    position(robot(R),(X,Y),T-1).
39 attr(node,(P,T),label,"R{{robot}}") :- position(robot(R),P,T), not position(shelf(_),P,T).
40 attr(node,(P,T),label,"S{{shelf}}") :- not position(robot(_),P,T), position(shelf(S),P,T).
41 attr(node,(P,T),label,"R{{robot}}S{{shelf}}") :- position(robot(R),P,T),
42    position(shelf(S),P,T).
43 attr(node,(P,T),(label,robot),R) :- position(robot(R),P,T).
44 attr(node,(P,T),(label,shelf),S) :- position(shelf(S),P,T).
47 attr(node,(P,T),shape,"point") :- free(P,T).
50 attr(node,(P,T),shape,"circle") :-occo(P,T,robot(_)).
53 attr(node,(P,T),color,white) :- free(P,T).
59 attr(node,(P,T),colorscheme,"blues9") :-occo(P,T,robot(_)).
60 attr(node,(P,T),fillcolor,R) :-occo(P,T,robot(R)).

```

```

clingo asprilo.lp instance.lp -c horizon=19 --outf=2 | \
clingraph --viz-encoding=viz-asprilo.lp --engine=neato | \
--out=animate --sort=asc-int | \
--select-model=0 --type=digraph

```

The initial call to *clingo* takes the problem encoding and instance and yields a plan of length 19, executed on a 7×7 grid with three robots, three shelves, and one picking station. The individual 20 images underlying the resulting animation are given in Figure 3. In these images, robots are represented by solid blue circles, shelves by solid orange squares, and the only picking station by a solid green circle. The initial frame of the animation shows all three robots positioned in the lower left corner. Subsequent images depict the robots' movements, indicated by black arrows, as they approach their respective shelves. In the fifth image, robot 3 reaches shelf 1, prompting an update of the labels to signify their shared location. The act of picking up the shelf is visually represented in the sixth image by a thicker border added around the shelf. Succeeding steps will show how both robot and shelf move together until they reach the picking station. The movements of the remaining robots unfold in a similar fashion.

Let us explain how this works in the visualization encoding by focusing on unoccupied nodes and robots; shelves and picking stations are treated analogously. An unoccupied position p at a time step t is captured by `free(p,t)` in Line 10. Similarly, `occo($p,t,robot(r)$)` tells us that robot r is the only object on position p at a time step t . This is thus neither derivable when a robot is under a shelf, carrying one, or at a picking station. With this in mind, we see that Lines 47 and 53 depict a position as a circle on a white node (plus omitted details) whenever the position is free. And, analogously Lines 50, 59, and 60 represent solitary robots by solid blue circles. Here, robots are differentiated using multiple shades of blue via the *graphviz* attribute `colorscheme`, where each robot selects one color option using an integer in attribute `fillcolor`. Once a robot shares a position with a shelf or picking station, the graphical representation changes (and instead the robot adopts the one of the shelf or picking station).

Moreover, a robot's label changes whenever it is under a shelf or carries one. This is handled in Lines 39–44. A template (cf. Section 4) is selected depending on whether there is a robot, a shelf, or both in the corresponding position (Lines 39, 40, and 41 to 42, respectively). The variables `robot` and `shelf`, used in the templates, are defined in Lines 43 and 44, respectively.

3.3 Scalable Vector Graphics and interactivity

An image format of particular interest is the Scalable Vector Graphics (SVG)¹⁷ format as it supports interactivity. More precisely, SVG is a text-based web standard for describing images in XML format integrated with Cascading Style Sheets (CSS) and JavaScript. In order to allow for interactive actions on graphic elements, we extend the SVG capabilities supported by *graphviz*. Our extension is implemented in JavaScript; it listens to events being performed on an element and reacts by changing a CSS style property on another element. To this end, the `class` attribute assigned to an element defines how the element changes on a given event: `click`, `mouseenter`, `mouseleave`, and `contextmenu`

¹⁷ <https://www.w3.org/TR/SVG/Overview.html>.

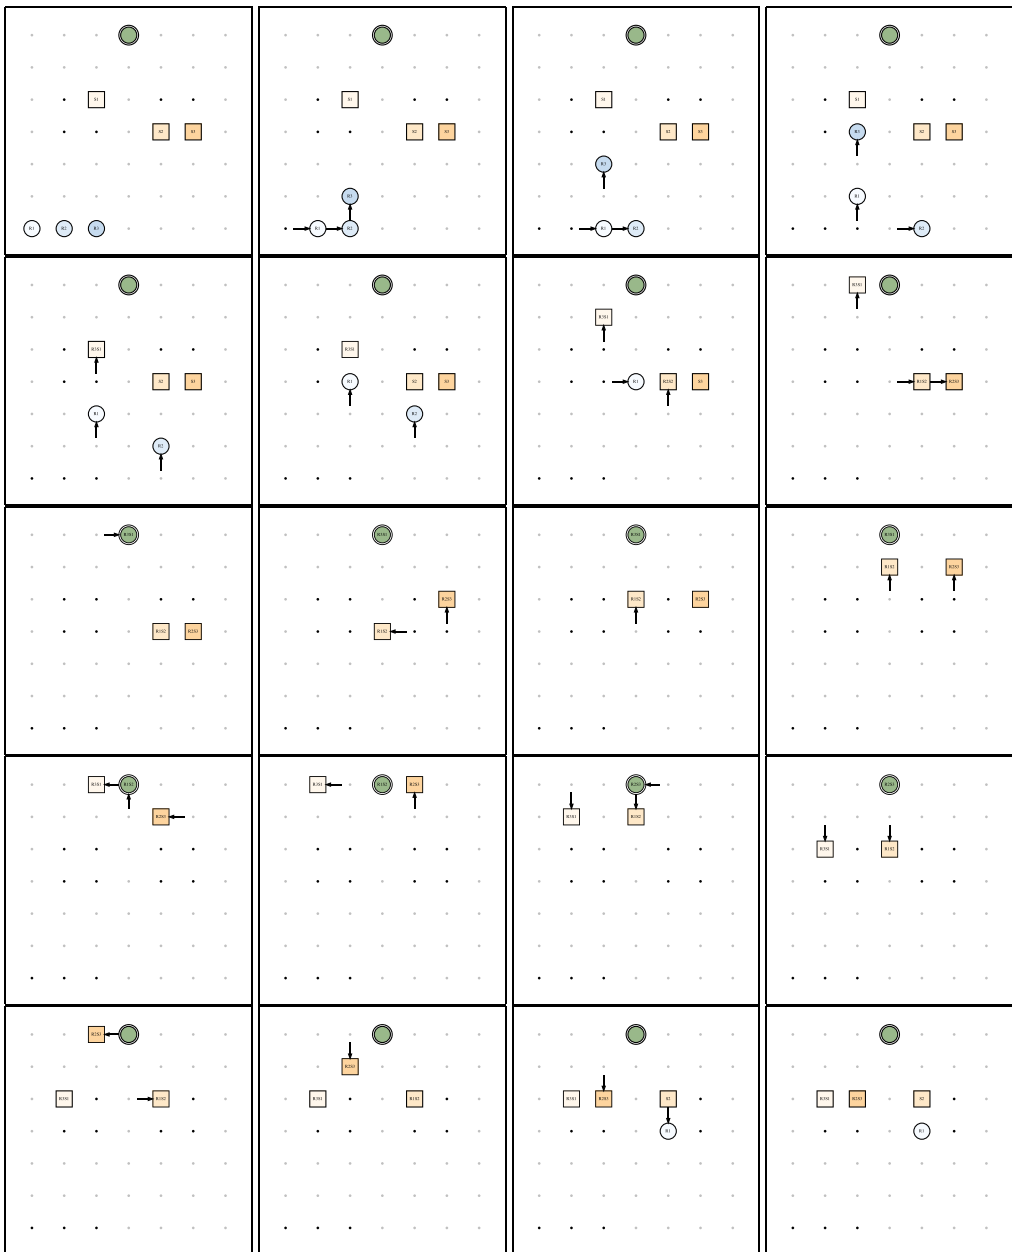


Fig. 3. Individual graph representations making up an animated plan.

(right click). Notably, all interaction is single-shot. That is, the resulting SVG file is generated once by a single call to *clingo* and no further interplay with the solver is possible. Therefore, all information for interactivity needs to be rendered in the same SVG file and no information of what actions are taken can be returned to the solver.

For illustration, we visualize the mouse-driven expansion of simple trees, defined by predicates `node/1`, `parent/2`, and `root/1` (see Listing 5 for an example). The

```

1  attr(node,X,style,"filled") :- node(X).
2  edge((X,Y)) :- parent(X,Y).

4  react(node,X,X) :- node(X).
5  react(node,Y,X) :- edge((X,Y)).
6  react(edge,(X,Y),X) :- edge((X,Y)).

8  attr(node,X,class,@svg_init("visibility","hidden")) :- node(X), not root(X).
9  attr(edge,E,class,@svg_init("visibility","hidden")) :- edge(E).
10 attr(T,E,class,@svg("click",X,"visibility","visible")) :- react(T,E,X).
11 attr(T,E,class,@svg("mouseenter",X,"opacity","1")) :- react(T,E,X).
12 attr(T,E,class,@svg("mouseleave",X,"opacity","0.2")) :- react(T,E,X).

```

Listing 4. *Specification of interaction on trees based on clingraph's SVG capabilities (tree-viz.lp)*

```

1  root(a).
2  node(a). node(b). node(c). node(d). node(e). node(f).
3  parent(a,b). parent(a,d). parent(b,c). parent(b,e). parent(d,f).

```

Listing 5. *Facts representing a simple tree (mytree.lp)*

corresponding visualization encoding is given in Listing 4. Line 1 adds a filled style to the nodes and Line 2 generates an edge for each instance of predicate `parent/2`. Lines 4–6 assign the reactions based on the underlying graph element using atoms of form `react(t,e1,e2)` that are read as: “element *e*₁ of type *t* reacts to actions on element *e*₂”. These atoms affect the visibility and opacity of nodes. For instance, Line 5 tells us that if there is an edge from *X* to *Y* then node *Y* reacts to actions on *X*. Lines 8–12 define the interactivity of the elements by assigning their `class` attribute to a formatted string. This string is handled by our extension, while the formatting of the strings is done by functions `@svg_init(property,value)` and `@svg(event,element,property,value)`. For instance, in Lines 8 and 9, function `@svg_init` is used to express that node *X* and edge *E* have the initial value `hidden` for property `visibility`. Line 10 states that an element *E* changes the value of the CSS property `visibility` to `visible` when *X* is clicked. The function `@svg` generates the string `clicked__X__visibility__visible` which is assigned as a `class` of *E* in the SVG file. The string is then parsed by our extension and mapped into the JavaScript method `addEventListener` to react when *X* is clicked. Similarly, Lines 11 and 12 define that an element *E* changes the value of `opacity` to 1 or 0.2, whenever the mouse enters or leaves element *X*, respectively.

As an example, consider the simple tree represented as facts in Listing 5. Together with the visualization encoding in Listing 4, it can be turned into an interactive image in SVG format by means of the following command:

```
clingraph mytree.lp --viz-encoding=tree-viz.lp --out=render --format=svg
```

A possible user interaction via mouse actions is indicated in Figure 4 via a series of snapshots.

While the simple black pointer highlights positions of interest, the one with surrounded by lines indicates a previous click. Each such click leads to the expansion of the tree by the succeeding nodes. Note how the opacity of a node and its subnodes changes whenever the pointer hovers over and away from it, respectively.

The transformation of *graphviz* into SVG uses the group element `<g>` to group all elements related to a node or edge. Since only these group elements can be indexed in JavaScript, the CSS style properties are set on the SVG group. This results in the

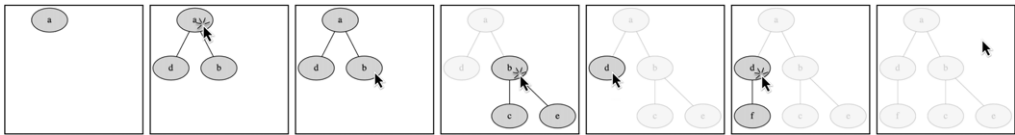


Fig. 4. Example user interaction via mouse actions expanding the tree.

```

1 attr(node,(X,Y),fillcolor,@svg_color()) :- cell(X,Y).
3 attr(node,(X,Y),class,@svg_init("color","gray")) :- cell(X,Y), (X+Y)\2 != 0.
4 attr(node,(X,Y),class,@svg_init("color","white")) :- cell(X,Y), (X+Y)\2 == 0.
5 attr(node,C,class,@svg("mouseenter",Q,"color","red")) :- attack(Q,C).
6 attr(node,(X,Y),class,@svg("mouseleave",Q,"color","gray")) :- attack(Q,(X,Y)),
7   (X+Y)\2 != 0.
8 attr(node,(X,Y),class,@svg("mouseleave",Q,"color","white")) :- attack(Q,(X,Y)),
9   (X+Y)\2 == 0.

```

Listing 6. *SVG interactive queens (svg-queens.lp)*

limitation that CSS style properties are not overwritten on the elements contained in the group. Thus, many property changes have no impact. In particular, this issue leads to problems when changing the color of elements. To address this issue, we offer a workaround for changing colors dynamically. That is, we provide the function `@svg_color` to represent the CSS value `currentcolor`. This can be used with any *graphviz* color attribute, such as `color`, `fillcolor`, and `fontcolor`, and serves as a placeholder for the color set using the SVG class.

For illustrating this functionality, Listing 6 extends the queens example with interactive elements for visualizing all cells attacked by a queen once it is hovered over by the mouse pointer. To this end, we replace Lines 6 and 7 in the visualization encoding in Listing 2 by Line 1 in Listing 6 to use the color set by the interactions specified in the following lines. Lines 3 and 4 set the initial color of the nodes to `gray` and `white`, respectively. Line 5 adds functionality to a node `C` attacked by queen `Q` so that the color is set to `red` when the mouse enters `Q`. When the mouse leaves `Q`, Lines 7 and 9 set the color back to the original value. We illustrate this by means of three snapshots in Figure 5. In the one on the left, the mouse pointer hovers over position (2,2). Accordingly, all cells on the same row, column, and diagonals are colored in red. The same happens in the two other scenarios, though initiated from position (5,1) and (1,4), respectively.

The approach has further limitations. For example, labels are independent of a CSS style and thus cannot be changed interactively. A way around this is to create multiple layers of nodes with the same position and change their visibility.¹⁸ However, we have no control over which elements appear on top and which on the bottom. Rather this must be handled manually by assuring that only a single element is visible in each position at each time. Another issue is that the position of all elements is fixed; therefore, expanding the size of the image on demand is impossible, only its visibility can be changed.

3.4 Visualizing the solving process of a Sudoku puzzle

Up to now, all case studies take answer sets as input for visualization. For the next example, however, we visualize partial assignments appearing during the search process

¹⁸ A nice example for this is the *minesweeper* puzzle given at <https://github.com/potassco/clingraph/tree/master/examples/minesweeper>.

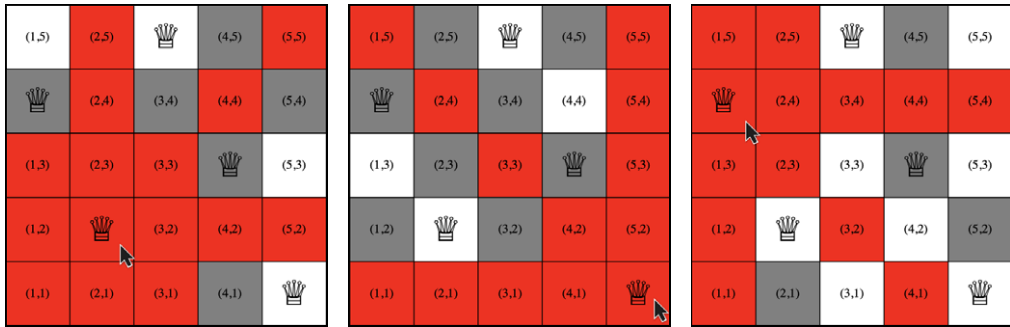


Fig. 5. Three snapshots of the user hovering the mouse pointer over specific cells on the board.

of *clingo*. Specifically, we discuss a visualization of the solving process of a Sudoku puzzle. To this end, we rely on *clingo*'s capability of integrating user-defined propagators¹⁹ into the solving process and use *clingraph*'s API for streamlining the declarative visualization of partial assignments.

In Table 2, we provide a generic propagator that can be used directly to monitor solving or as a template to create a domain-specific propagator. Basically, `ClingraphPropagator` class implements the user-defined propagator interface expected by *clingo*'s Python API. Its instance variables are initialized in Lines 10–12. The `viz_encoding` variable holds the path of the visualization encoding specific to the problem domain. The propagator uses this encoding to generate the facts defining the graph to visualize each partial assignment, which are stored in the `factbases` list. Additionally, the instance variable `l2s` maps each literal used internally by *clingo* to the corresponding list of atomic symbols from the problem encoding. Specifically, this mapping is formed in Line 17, just before solving process starts, when *clingo* calls the `init` function (Line 14) of `ClingraphPropagator`. Note that in Line 18 the propagator requests the solver to be notified when the truth value of these internal literals changes. Hence, with the help of `l2s`, the propagator functions can find the corresponding atoms of a solver literal whose truth value has changed during solving.

The main functionality of the propagator is to compile and prepare partial assignments appearing during various stages of the search process as reified atoms, which are passed to the visualization encoding. Such facts are of the form `_true(a)`, `_false(a)` and `_undefined(a)` for each atom *a* if it is assigned to true, false or neither in the current partial assignment, respectively. The key stages account for times when *clingo* reaches a fixpoint during unit propagation; decides on a literal; or faces a conflict and is about to backtrack. In each situation, *clingo* calls the corresponding propagator function `propagate` (Line 20), `decide` (Line 51) or `undo` (Line 35), respectively, and makes the partial assignment accessible to them. Hence, these functions are suitable for preparing the reified atoms of the partial assignment at the time of the call. In the `propagate` function, for instance, these facts are generated in Lines 23–27 and functions `decide` and `undo` have the same corresponding statements. Note that for each solver literal, corresponding atoms are found via the mapping `l2s` and truth values of such atoms are queried in

¹⁹ <https://potassco.org/clingo/python-api/current/clingo/propagator.html>.

Table 2. *The propagator class for visualizing solving*

```

8 class ClingraphPropagator:
9     def __init__(self, viz_encoding):
10         self.viz_encoding = viz_encoding
11         self.factbases = []
12         self.l2s = {}

14     def init(self, init): ...
15         for atom in init.symbolic_atoms:
16             lit = init.solver.literal(atom.literal)
17             self.l2s.setdefault(lit, []).append(str(atom.symbol))
18             init.add_watch(lit)

20     def propagate(self, ctl, changes):
21         i = len(self.factbases)
22         propagation_prg = [f"_step_type(propagate, {i}).", f"_level({ctl.assignment.decision_level})."]
23         for l, symbols in self.l2s.items():
24             v = ctl.assignment.value(l)
25             t = '_undefined' if v is None else '_true' if v else '_false'
26             for s in symbols:
27                 propagation_prg.append(f"{t}({s}). ")
28         for l in changes:
29             symbols = self.l2s[l]
30             for s in symbols:
31                 propagation_prg.append(f"_change({s}). ")
32         self.add_factbase(propagation_prg)
33         return True

35     def undo(self, solver_id, assign, undo):
36         i = len(self.factbases)
37         propagation_prg = [f"_step_type(undo, {i}).", f"_level({assign.decision_level})."]

39         for l, symbols in self.l2s.items():
40             v = assign.value(l)
41             t = '_undefined' if v is None else '_true' if v else '_false'
42             for s in symbols:
43                 propagation_prg.append(f"{t}({s}). ")
44         for l in undo:
45             symbols = self.l2s[l]
46             for s in symbols:
47                 propagation_prg.append(f"_change({s}). ")

49         self.add_factbase(propagation_prg)

```

Table 2. *Continued*

```

51     def decide(self, thread_id, assign, fallback):
52         i = len(self.factbases)
53         propagation_prg = [f"_step_type(decide,{i}).", f"_level({assign.decision_level})."]

55         for l,symbols in self.l2s.items():
56             v = assign.value(l)
57             t = '_undefined' if v is None else '_true' if v else '_false'
58             for s in symbols:
59                 propagation_prg.append(f"{t}({s}).")
60         if abs(fallback) in self.l2s:
61             for s in self.l2s[abs(fallback)]:
62                 pol = "pos" if fallback > 0 else "neg"
63                 propagation_prg.append(f"_decide({s},{pol}).")

65         self.add_factbase(propagation_prg)
66         return 0

68     def add_factbase(self, prg_list):
69         fb = Factbase()
70         ctl = Control([])
71         ctl.load(self.viz_encoding)
72         ctl.add("base", [], "".join(prg_list))
73         ctl.ground([( "base", [])],ClingraphContext())
74         ctl.solve(on_model=fb.add_model)
75         self.factbases.append(fb)

```

the current partial assignment in Line 24. Additionally, in each stage we generate the fact `_step_type(t,i)` where *t* is either `propagate`, `decide` or `undo`, and *i* is a natural number identifying the solving step (in Lines 22, 37, and 53). Such facts are required not only to designate the type of the current stage but also to order the visualization of each generated partial assignment. This ordering allows us to represent *clingo*'s solving process by combining individual graphs as an animation. The functions `propagate` and `undo` generate additional facts of the form `_change(a)`, where the truth value assignment to atom *a* has changed during the propagation (Lines 28–31) or will be undone during backtracking (Lines 44–47), respectively. Similarly, the function `decide` generates a reified fact of the form `_decide(a,p)` in Lines 60–63 to represent *clingo* chooses atom *a* to have a truth value of true or false depending on *p* being `pos` or `neg`, at a decision point, respectively.

For each solving stage, we process the reified atoms of the active partial assignment with the problem domain's visualization encoding to generate the facts defining the graph. This is achieved by calling the `add_factbase` function defined in Lines 68–75 at the end of each solving stage. Each resulting graph facts gets stored in a `Factbase` object of *clingraph*'s API in Line 75. Once *clingo*'s solving is done, we process all `Factbase` objects accumulated in the propagator using *clingraph* to generate individual graphs for each of the partial assignments. Finally, we combine these graphs to generate an animation of *clingo*'s solving process. Unlike the previous examples, we rely on *clingraph*'s API functions (eg., `compute_graphs` and `save_gif`) to carry out these tasks.

To illustrate the process described above, we use the Sudoku puzzle from *clingraph*'s examples folder.²⁰ In this encoding, we use predicate `sudoku(x,y,v)` to represent a cell with coordinates (*x*,*y*) in a 9 × 9 grid with an assigned digit *v* from 1 to 9. A cell can have an initial value defined in the instance by predicate `initial(x,y,v)` or it can be empty if no such predicate appears. Then, the problem encoding and instance are handed to *clingo*'s solving process which is observed by our propagator. Partial assignments accumulated by the propagator are passed to the visualization encoding, which is shown in Table 3. Additionally, Figure 6 depicts the resulting animation's key frames visualizing the partial assignments reached during solving.

Let us now examine how the frames from Figure 6 are constructed. Each cell with an initial value is visualized by setting the corresponding digit as the label of its node (rule in Line 39 from Table 3) and using a relatively larger font size (rule in Line 38). These rules have the reified literal `_true(initial(X,Y,V))` in the body to represent cells with initial values. Notice that facts appearing in the problem input, such as `initial(X,Y,V)`, will always have their truth value set to true. For each node of an empty cell, we construct an HTML-like label that allows us to use rich visual elements like tables with different borders and background colors. In order to ease constructing long HTML-like labels, we rely on template strings (see Section 4). Let us first cover empty cells that must be filled with one specific digit. The HTML-like label for such a cell represents a table having only one slot for the respective digit. The rule in Lines 43–48 generates such a label as a template string by concatenating the constituent strings using the `concat` external function provided by *clingraph*. Note that the rule body designates an initially empty cell (captured by the body literal `not _true(initial(X,Y,_))`) that must be filled with a

²⁰ <https://github.com/potassco/clingraph/tree/master/examples/propagator/sudoku>.

Table 3. Selected lines from the encoding visualizing Sudoku solving (*viz-sudoku-solving.lp*)

```

38 attr(node, pos(X,Y), fontsize,40 ) :- _true(initial(X,Y,V)).
39 attr(node, pos(X,Y), label, V ) :- _true(initial(X,Y,V)).

42 attr(node, pos(X,Y), fontsize, 20) :- _true(sudoku(X,Y,V)), not _true(initial(X,Y,_)).
43 attr(node, pos(X,Y), label, @concat("<<table BORDER='0'>",
44   "<tr><td BGCOLOR='{{color[",V,"]}}'{{opacity[",V,"]}}'>",
45   "  BORDER='{{border[",V,"]}}'>",
46   "  {{value[",V,"]}}",
47   "</td></tr></table>>")) :-
48   _true(sudoku(X,Y,V)), not _true(initial(X,Y,_)).

52 attr(node, pos(X,Y), label, @concat("<<table BORDER='0'>",
53   "<tr>",
54   "<td BGCOLOR='{{color[1]}}'{{opacity[1]}}' BORDER='{{border[1]}}'>{{value[1]}}</td>",
55   "<td BGCOLOR='{{color[2]}}'{{opacity[2]}}' BORDER='{{border[2]}}'>{{value[2]}}</td>",
56   "<td BGCOLOR='{{color[3]}}'{{opacity[3]}}' BORDER='{{border[3]}}'>{{value[3]}}</td>",
57   "</tr>",
58   "<tr>",
59   "<td BGCOLOR='{{color[4]}}'{{opacity[4]}}' BORDER='{{border[4]}}'>{{value[4]}}</td>",
60   "<td BGCOLOR='{{color[5]}}'{{opacity[5]}}' BORDER='{{border[5]}}'>{{value[5]}}</td>",
61   "<td BGCOLOR='{{color[6]}}'{{opacity[6]}}' BORDER='{{border[6]}}'>{{value[6]}}</td>",
62   "</tr>",
63   "<tr>",
64   "<td BGCOLOR='{{color[7]}}'{{opacity[7]}}' BORDER='{{border[7]}}'>{{value[7]}}</td>",
65   "<td BGCOLOR='{{color[8]}}'{{opacity[8]}}' BORDER='{{border[8]}}'>{{value[8]}}</td>",
66   "<td BGCOLOR='{{color[9]}}'{{opacity[9]}}' BORDER='{{border[9]}}'>{{value[9]}}</td>",
67   "</tr>",
68   "</table>>")) :-
69   _true(pos(X,Y)), not _true(initial(X,Y,_)), not _true(sudoku(X,Y,_)).

74 attr(node, pos(X,Y), (label,opacity,V), 25 ) :- not _change(sudoku(X,Y,V)),
75   not _decide(sudoku(X,Y,V),_), _true(sudoku(X,Y,V)).
76 attr(node, pos(X,Y), (label,opacity,V), "00") :- not _true(sudoku(X,Y,V)),
77   not _decide(sudoku(X,Y,V),neg), _true(pos(X,Y)), value(V).

80 attr(node, pos(X,Y), (label,border,V), 1) :- _decide(sudoku(X,Y,V),_).
81 attr(node, pos(X,Y), (label,border,V), 0) :- not _decide(sudoku(X,Y,V),_),
82   _true(pos(X,Y)), value(V).

85 attr(node, pos(X,Y), (label,color,V), white) :- not _true(sudoku(X,Y,V)),
86   not _decide(sudoku(X,Y,V),_), _true(pos(X,Y)), value(V).
87 attr(node, pos(X,Y), (label,color,V), green) :- _true(sudoku(X,Y,V)).
88 attr(node, pos(X,Y), (label,color,V), red ) :- _decide(sudoku(X,Y,V),neg).
89 attr(node, pos(X,Y), (label,color,V), green) :- _decide(sudoku(X,Y,V),pos).

```

Table 3. *Continued*

```
92 attr(node, pos(X,Y), (label,value,V), V ) :- _true(sudoku(X,Y,V)).
93 attr(node, pos(X,Y), (label,value,V), V ) :- _undefined(sudoku(X,Y,V)).
94 attr(node, pos(X,Y), (label,value,V), "" ) :- _false(sudoku(X,Y,V)).
```

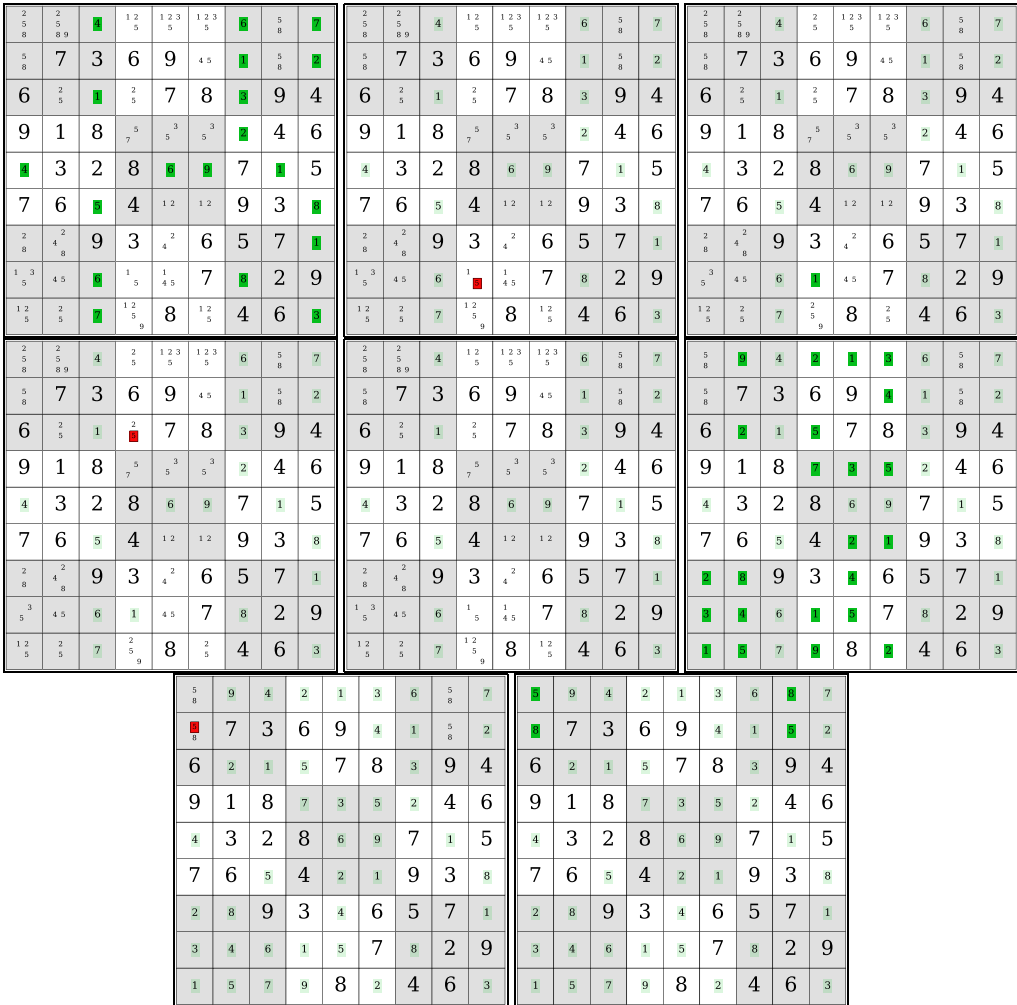


Fig. 6. Visualizations of the stages while solving a Sudoku puzzle.

specific digit (`_true(sudoku(X,Y,V))`). We set the font size of these cells via the rule in Line 42. An example of this can be found on the third cell in the topmost row of the top leftmost graph in Figure 6, where this initially empty cell is now filled with digit 4 with dark green background. For adding style to such cells, the label template uses variables to represent the color, opacity, border, and value. The values for these variables are obtained through different rules that generate atoms over `attr/4`. For this specific cell, the RGB code of dark green as the value of variable `color[V]` is set via the rule in Line 87. Furthermore, the rule in Lines 81–82 assigns value 0 to variable `border[V]` to avoid setting borders and the rule in Line 92 assigns 4 to `value[V]`. We also add opacity to the background color codes to highlight changes in the current partial assignment from the ones propagated in previous assignments by reducing the opacity of older ones. In this specific cell, since variable `opacity[V]` gets empty string as a default value due to absence of any rule generating a specific value for the variable, the opacity of the background color is not modified. The same cell but in the following graph, has a light green background

color, which designates that *clingo* filled it in an earlier propagation step. In order to generate the light green color, in that step variable `opacity[V]` gets its value 25 via the rule in Lines 74–75. Note that its body literal `not _change(sudoku(X,Y,V))` captures the respective cell is not filled in the current partial assignment. These rules generating variable values are also used for the Sudoku cells that have multiple options. We describe such cells below.

The remaining type of cells are those initially empty cells in which more than one digit may appear in a partial assignment. Their HTML-like label represents a 3×3 table allowing a slot for each digit from 1 to 9. Our aim is to visualize digits that can possibly appear in such a cell in this tabular form. For instance, the top leftmost graph shows that either 2, 5 or 8 can be placed in the first cell. The rule in Lines 52–69 constructs the label as a template string representing the 3×3 table. The body literal `not _true(sudoku(X,Y,_))` captures these initially empty cells with multiple options. Consider *clingo* has reasoned that digit d from 1 to 9 cannot appear in an empty cell (x,y) in a partial assignment, which is represented by the reified fact `_false(sudoku(x,y,d))` generated by the propagator (e.g., `_false(sudoku(1,9,1))` for the cell mentioned above). We do not show d in its respective slot. To this end, template variable `value[d]` is assigned to empty string via the rule in Line 94. Additionally, its transparent background color is controlled via rules in Lines 85–86 and in Lines 76–77 by setting `white` to variable `color[d]` and "00" to variable `opacity[d]`, respectively. Also consider *clingo* may be undecided on whether digit d is the value of the empty cell or not (e.g., digits 2, 5, and 8 for the cell mentioned above). This is reflected by the fact `_undefined(sudoku(x,y,d))` generated by the propagator for a partial assignment. We show d in its respective slot by setting variable `value[d]` to d this time via the rule in Line 93. Its transparent background is set via the same rules in Lines 85–86 and Lines 76–77. We can also visualize whenever the propagation during solving reaches a fixpoint, and *clingo* may decide on a truth value of an undefined atom to continue search. For instance, the second graph in the first row of Figure 6 shows such a decision point as digit 5 in red background with a border where *clingo* selects the atom `sudoku(4,2,5)` to be false. Its background color and border are set via rules in Lines 88 and 80, respectively. Whenever *clingo* selects an atom to be true at a decision point, we visualize it as green (rule in Line 89).

Ultimately, our animation allows us to analyze different aspects of the solving process of the Sudoku. For instance, the first graph illustrates that during the initial propagation *clingo* already fills many cells with digits (those having digits with green background) and constrains the remaining empty cells that only possible digits are shown. This can be an indicator of how simple the Sudoku instance is. Finally, when we reach the last graph (bottom rightmost) passing through various stages of solving in order, we get an answer set representing a solution of the puzzle instance.

3.5 Visualizing the program structure

So far, we have visualized the solving process, input, and/or result of a program. However, we may also visualize information about the program itself. In this section, we concentrate on the abstract syntax tree (AST) of a program, as it is accessible via the `clingo.ast` module of *clingo*'s Python API. Visualizing a program's AST eases the understanding of its internal structure. This is of particular interest when dealing with non-ground

```

1 ast_node(875,type,"AST").
2 ast_node(875,variant,"Rule").
3 ast_node(875,value,"#false :- edge(N,M), assign(N,C),
4 assign(M,C).").
5 ast_edge((875,885),key,"head").
6 ast_edge((875,898),key,"body").

```

Listing 7. Partial translation of the program in Listing 1 into our intermediate fact format

programs. To this end, we follow a three-stage process. First, we translate the AST into an intermediate fact format using a simple Python script called `reify_ast.py`. Then, we employ a visualization encoding to convert these facts into *clingraph*'s format. Finally, we run *clingraph* to render the graph. All examples can be found in *clingraph*'s repository.²¹

Our intermediate fact format uses two predicates: `ast_node/3` and `ast_edge/3` represent the nodes and edges in an AST. These predicates are meant to be semantic triples, linking a subject, in our case a node or edge, via a key to a value. This is inspired by the representation of graph data, as used in the Resource Description Framework²². To illustrate how ASTs are stored in terms of triples, let us look at the translation of the program in Listing 1. It is partially depicted in Listing 7 and may be obtained by running `reify_ast.py color.lp`.

Nodes are represented by unique integers which are arbitrarily chosen by our script. Edges are identified using the integers of the adjacent nodes. Line 1 tells us that a node 875 exists and that it stands for an instance of the class `clingo.ast.AST` in the Python API of *clingo*. The other possible types are `ASTSequence`, `Location`, `Position`, `Symbol`, `int`, `str` and `None`. They cover the respective classes in `clingo.ast` and the necessary basic types in Python. If the type of a node is `AST`, declaring a variant as in Line 2 is mandatory. The variant reflects the `clingo.ast.ASTType` of each instance of `clingo.ast.AST`. The possible variants include `Rule`, `Variable`, `SymbolicAtom`, and many more. Line 4 assigns a value to node 875, in this case a string representing the rule represented by node 875. Lines 5 and 6 reflect two outgoing edges of node 875, named `head` and `body`. As the names suggest, these edges point to nodes capturing the head and body of the rule.

In order to translate the output of our script into the input format of *clingraph*, we employ a visualization encoding, assembling an HTML-like label including all the data stored in the semantic triples. Its main component is the template (see Section 4) shown in Listing 8 from Lines 10 to 23. Given that even small programs have large syntax trees, our encoding provides functionalities (Lines 1–5) to show and hide subtrees. For instance, using `ast_show(node, 875)` guarantees that only the subtree of node 875 is shown, while the rule

```

ast_hide(edge, I) :- ast_edge(I, _, _), I = (_, I2),
                    ast_node(I2, type, "Location").

```

allows us to hide any subtree that is rooted at a node of type `Location`. Calling

```

reify_ast.py color.lp | clingraph --viz-encoding=viz_ast.lp \
                                --type=digraph \
                                --out=render --format=pdf

```

²¹ <https://github.com/potassco/clingraph/tree/master/examples/ast>.

²² <https://www.w3.org/TR/rdf11-concepts/>.

```

1  ast_show(edge, I) :- ast_edge(I, _, _), I = (I1, _), ast_show(node, I1).
2  ast_show(node, I) :- ast_show(edge, (_, I)).

4  ast_hide(node, I) :- ast_hide(edge, (_, I)).
5  ast_hide(edge, I) :- ast_edge(I, _, _), I = (I1, _), ast_hide(node, I1).

7  node(I) :- ast_node(I, _, _), ast_show(node, I), not ast_hide(node, I).
8  edge(I) :- ast_edge(I, _, _), ast_show(edge, I), not ast_hide(edge, I).

10 attr(node, I, label, @concat(
11     "<<table border='0' cellborder='1' cellspacing='0' cellpadding='3'>",
12     "<tr>",
13         "<td>{{id}}</td>",
14         "<td>{{type}}</td>",
15         "{% if variant %}<td>{{variant}}</td>{% endif %}",
16     "</tr>",
17     "{% if value %}",
18         "<tr><td colspan='{{colspan}}'>",
19             "<font face='monospace'> {{value}} </font>",
20         "</td></tr>",
21     "{% endif %}",
22     "</table>>"
23 )) :- node(I).

25 attr(node, I, (label, id), I)
26     :- node(I).
27 attr(node, I, (label, type), T)
28     :- node(I), ast_node(I, type, T).
29 attr(node, I, (label, variant), V)
30     :- node(I), ast_node(I, variant, V).
31 attr(node, I, (label, colspan), 2)
32     :- node(I), ast_node(I, value, _), not ast_node(I, variant, _).
33 attr(node, I, (label, colspan), 3)
34     :- node(I), ast_node(I, value, _), ast_node(I, variant, _).
35 attr(node, I, (label, value), @html_escape(V))
36     :- node(I), ast_node(I, value, V).

38 attr(edge, I, label, L) :- edge(I), ast_edge(I, key, L).

40 attr(graph_nodes, default, fontsize, 10).
41 attr(graph_nodes, default, shape, plain).

43 attr(graph_edges, default, fontsize, 10).

```

Listing 8. Selected lines from the encoding visualizing the AST (*viz-ast.lp*)

with the above line in the visualization encoding instructs *clingraph* to render the graph shown in Figure 7.

Showing the AST is by far not the only option to visualize a program's structure. In principle, *clingraph* may render any structured knowledge about the program provided that a reification format, a tool generating it, and a visualization encoding exists. To this end, our case study may serve as a blueprint for future ideas.

4 Formatting attributes with templates

Generating complex string values for attributes can become quite cumbersome, especially when dealing with HTML-like labels.²³ This type of *graphviz* labels are formed by an

²³ <https://graphviz.org/doc/info/shapes.html#html>.

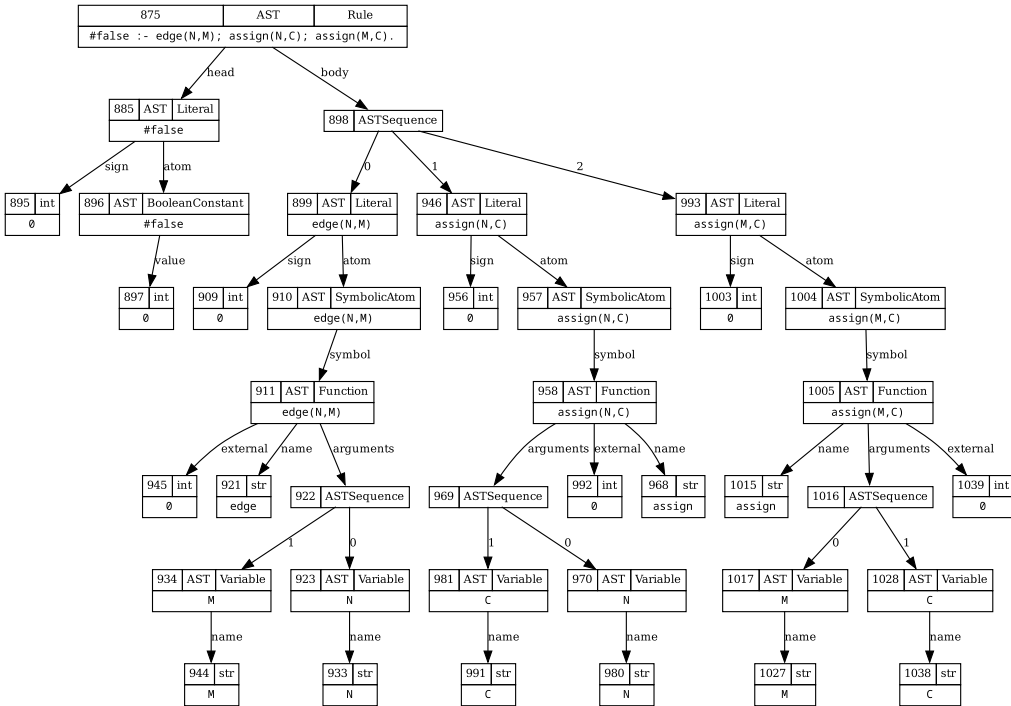


Fig. 7. A partial visualization of the AST of the program in Listing 1.

HTML string delimited by `<...>` which gives a lot of flexibility for formatting the text and generating tables. We simplify the generation of such strings by using the template engine Jinja.²⁴ Attribute values can then be seen as Jinja templates, which are rendered using the variables provided by atoms of the form `attr(t, id, (n, v), x)`. In these atoms, the third argument is a pair indicating that variable `v` has value `x` when rendering the template of attribute `n`. Furthermore, we can encapsulate values in dictionary variables by using triples instead, where `(n, v, k)` indicates that variable `v` is a dictionary with entry `{k : x}`. When no template is provided, the values of all the variables are concatenated.

For illustration, we visualize the data of people, defined by predicates `person/1`, `name/1`, `middlename/1` and `lastname/1` (see Listing 10 for an example). We visualize the data using HTML-like labels to generate the tables in Listing 9. Line 1 generates a node for each person and Line 2 removes the shape of the node (no shape is needed since the label is a table). Lines 3–7 define the template for the label using the HTML tags `<table>`, `<tr>`, `<td>` and `` to construct a table, row, cell, and boldface text, respectively. Variables are enclosed in double braces `{{...}}` and corresponding values are substituted by Jinja for these variables. The rules in Lines 8–11, for instance, generate atoms of `attr/4` to populate values for the template variables. Lines 8 and 11 use a pair to assign `N` to the variable `id` and the last name to `lastname`. Unlike, Lines 9 and 10 use a triple, making the variable `name` a dictionary with the keys `first` and `middle`, which is accessed in the template as `name['first']` and `name['middle']`, respectively. The output of

²⁴ <https://jinja.palletsprojects.com/en/3.1.x/templates>.


```

1 node(N) :- person(N).
2 attr(node,N,shape,none) :- person(N).
3 attr(node,N,label,@concat(
4   "<<table>",
5   "<tr><td><b>{{id}}</b></td></tr>",
6   "<tr><td>{{lastname}} ({{name['first']}} {{name['middle']}})</td></tr>",
7   "</table>>") :- person(N).
8 attr(node,N,(label,id),N) :- person(N).
9 attr(node,N,(label,name,first),Name) :- name(N,Name).
10 attr(node,N,(label,name,middle),Name) :- middlename(N,Name).
11 attr(node,N,(label,lastname),Lastname) :- lastname(N,Lastname).

```

Listing 9. Visualization encoding to exemplify the generation of strings via templates (*template.lp*).

```

1 person(anna).
2 name(anna,"Anna"). middlename(anna,"Julia"). lastname(anna,"Scott").
3 person(tom).
4 name(tom,"Thomas"). lastname(tom,"Blake").

```

Listing 10. Instance for the template example (*people.lp*).

tom	anna
Blake (Thomas)	Scott (Anna Julia)

Fig. 8. Example of HTML-like labels using attribute templates.

this encoding together with the instance defined in Listing 10 is shown in Figure 8. It is produced by means of the following instruction:

```
clingraph people.lp --viz-encoding=template.lp --out=render
```

Jinja's syntax of templates also includes statements like conditionals, loops, and several operations. We refer the interested reader to our github repository for more complex examples of *clingraph* using such features.²⁵

5 Related work

Many aspects of *clingraph* are inspired by previous systems described in the literature. The basic goal – to visualize answer sets by mapping special atoms to graphic elements – traces back to *aspviz* (Cliffe *et al.* 2008), a command-line application written in Java using the Standard Widget Toolkit (SWT) for rendering. It is capable of rendering two-dimensional graphics with absolute coordinates but does neither allow relative positioning nor graph structures. These features were introduced by *kara* (Kloimüller *et al.* 2011), a plugin written for the SeaLion IDE. The alternative of using *graphviz* as a backend was first mentioned by the authors of *aspviz* and followed up with a rather basic implementation in *lonsdaleite*.²⁶ Another visualizer for answer sets is *idpdraw*,²⁷ although it seems to be discontinued.

The idea of visualizing the solving process was first explored for the *nomore* system (Bösel *et al.* 2004) which uses a graph-oriented computational model. For *dlv*, there exists a graphical tool for developing and testing logic programs (Perri *et al.* 2007) as well

²⁵ <https://github.com/potassco/clingraph/tree/master/examples/office>.

²⁶ <https://github.com/rndmcnllly/Lonsdaleite>.

²⁷ <https://dtai.cs.kuleuven.be/krr/files/bib/manuals/IDPDraw-manual.pdf>.

as a visual tracer (Calimeri et al. 2009). In the realms of *clingo*, visualizing the solving process has been explored using a tweaked version of *clasp* (König and Schaub 2013).

Our system not only integrates ideas from the literature and makes them available for modern ASP systems, but also has some features that have – to the best of our knowledge – never been implemented before. There is a powerful API which makes it easy to include *clingraph* in custom projects, a multitude of different output formats including L^AT_EX and animated GIF, and the capacity of integrating a propagator for visualizing the solving process of *clingo*.

6 Discussion

Clingraph provides essentially an ASP-based front-end to the graph visualization software *graphviz*. In doing so, it takes up the early approach of *aspviz* (Cliffe et al. 2008) and extends it in the context of modern ASP technology. The advantage of *clingraph* is that one does not have to resort to foreign programming languages for visualization but rather remains within the realm of ASP. This provides users with an easy interface among logic programs and/or answer sets and their visualization. Moreover, *clingraph* offers a Python API that extends this ease of interfacing to *clingo*'s API, and in turn to connect and monitor various aspects of the solving process. The fact-based interface of *clingraph* makes it readily applicable to any ASP system. For more advanced features, like json output and API functionality, *clingraph* depends on *clingo*. *Clingraph* is open source software and freely available at <https://github.com/potassco/clingraph>.

Clingraph is a single-shot system mapping logic programs to graphs – no return flow is envisaged. A natural extension is thus to allow for returning user input into the system. This idea is pursued in the system *clinguin*²⁸ (Hahn 2023), which allows for interacting through user interfaces, specified in terms of logic programs. Another avenue of future research is to consider alternatives to graphs for visualization. Candidates of interest are plots and tables but also game descriptions seem possible.

Acknowledgments

This work was supported by DFG grants SCHA 550/11 and 15 as well as BMBF project ISCO with support code KK5291302GR1.

References

- BÖSEL, A., LINKE, T. AND SCHAUB, T. 2004. Profiling answer set programming: The visualization component of the noMoRe system. In *Proceedings of the Ninth European Conference on Logics in Artificial Intelligence (JELIA'04)*, J. Alferes and J. Leite, Eds. Lecture Notes in Computer Science, vol. 3229. Springer-Verlag, 702–705.
- CALIMERI, F., LEONE, N., RICCA, F. AND VELTRI, P. 2009. A visual tracer for DLV. In *Proceedings of the Second Workshop on Software Engineering for Answer Set Programming (SEA'09)*, M. de Vos and T. Schaub, Eds., vol. 546. CEUR Workshop Proceedings, 79–93.

²⁸ <https://github.com/potassco/clinguin>.

- CLIFFE, O., DE VOS, M., BRAIN, M. AND PADGET, J. 2008. ASPVIZ: Declarative visualisation and animation using answer set programming. In *Proceedings of the Twenty-fourth International Conference on Logic Programming (ICLP'08)*, M. Garcia de la Banda and E. Pontelli, Eds. Lecture Notes in Computer Science, vol. 5366. Springer-Verlag, 724–728.
- GEBSER, M., OBERMEIER, P., OTTO, T., SCHAUB, T., SABUNCU, O., NGUYEN, V. AND SON, T. 2018. Experimenting with robotic intra-logistics domains. *Theory and Practice of Logic Programming* 18, 3-4, 502–519.
- HAHN, S. 2023. clinguin: Building user interfaces in asp. In *Proceedings of the Sixth Workshop on Trends and Applications of Answer Set Programming (TAASP'23)*. To appear.
- HAHN, S., SABUNCU, O., SCHAUB, T. AND STOLZMANN, T. 2022. clingraph: Asp-based visualization. In *Proceedings of the Sixteenth International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR'22)*, G. Gottlob, D. Incezan and M. Maratea, Eds. Lecture Notes in Artificial Intelligence, vol. 13416. Springer-Verlag, 401–414.
- KLOIMÜLLNER, C., OETSCH, J., PÜHRER, J. AND TOMPITS, H. 2011. Kara: A system for visualising and visual editing of interpretations for answer-set programs. In *Proceedings of the Nineteenth International Conference on Applications of Declarative Programming and Knowledge Management (INAP'11)*, S. Abreu, J. Oetsch, J. Pührer, D. Seipel, H. Tompits, M. Umeda and A. Wolf, Eds. Number INFSYS RR-1843-11-06 in Institute for Information Systems, Technische Universität Wien, Technical Report Series, 325–344.
- KÖNIG, A. AND SCHAUB, T. 2013. Monitoring and visualizing answer set solving. In *Technical Communications of the Twenty-ninth International Conference on Logic Programming (ICLP'13)*, E. Lamma and T. Swift, Eds. Theory and Practice of Logic Programming, Online Supplement, vol. 13(4-5).
- LIFSCHITZ, V. 2019. *Answer Set Programming*. Springer-Verlag.
- PERRI, S., RICCA, F., TERRACINA, G., CIANNI, D. AND VELTRI, P. 2007. An integrated graphic tool for developing and testing dlv programs. In *Proceedings of the First Workshop on Software Engineering for Answer Set Programming (SEA'07)*, M. de Vos and T. Schaub, Eds., vol. 281. CEUR Workshop Proceedings, 86–100.