

Structural Change Distilling of Ansible Roles

Presentation Abstract

Ruben Opdebeeck

ropdebee@vub.be

Ahmed Zerouali

azeroual@vub.be

Camilo Velázquez-Rodríguez

cavelazq@vub.be

Coen De Roover

cderoove@vub.be

Ansible Galaxy





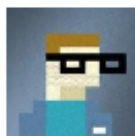



🔍 Search

nginx 🔍 ^ Filters (1458 results)

Type ▾ Filter by Collection or Role... ▾ | Best Match ▾ ↓ 1

1458 Results Active filters: Deprecated: False × Type: Role × [Clear All Filters](#)

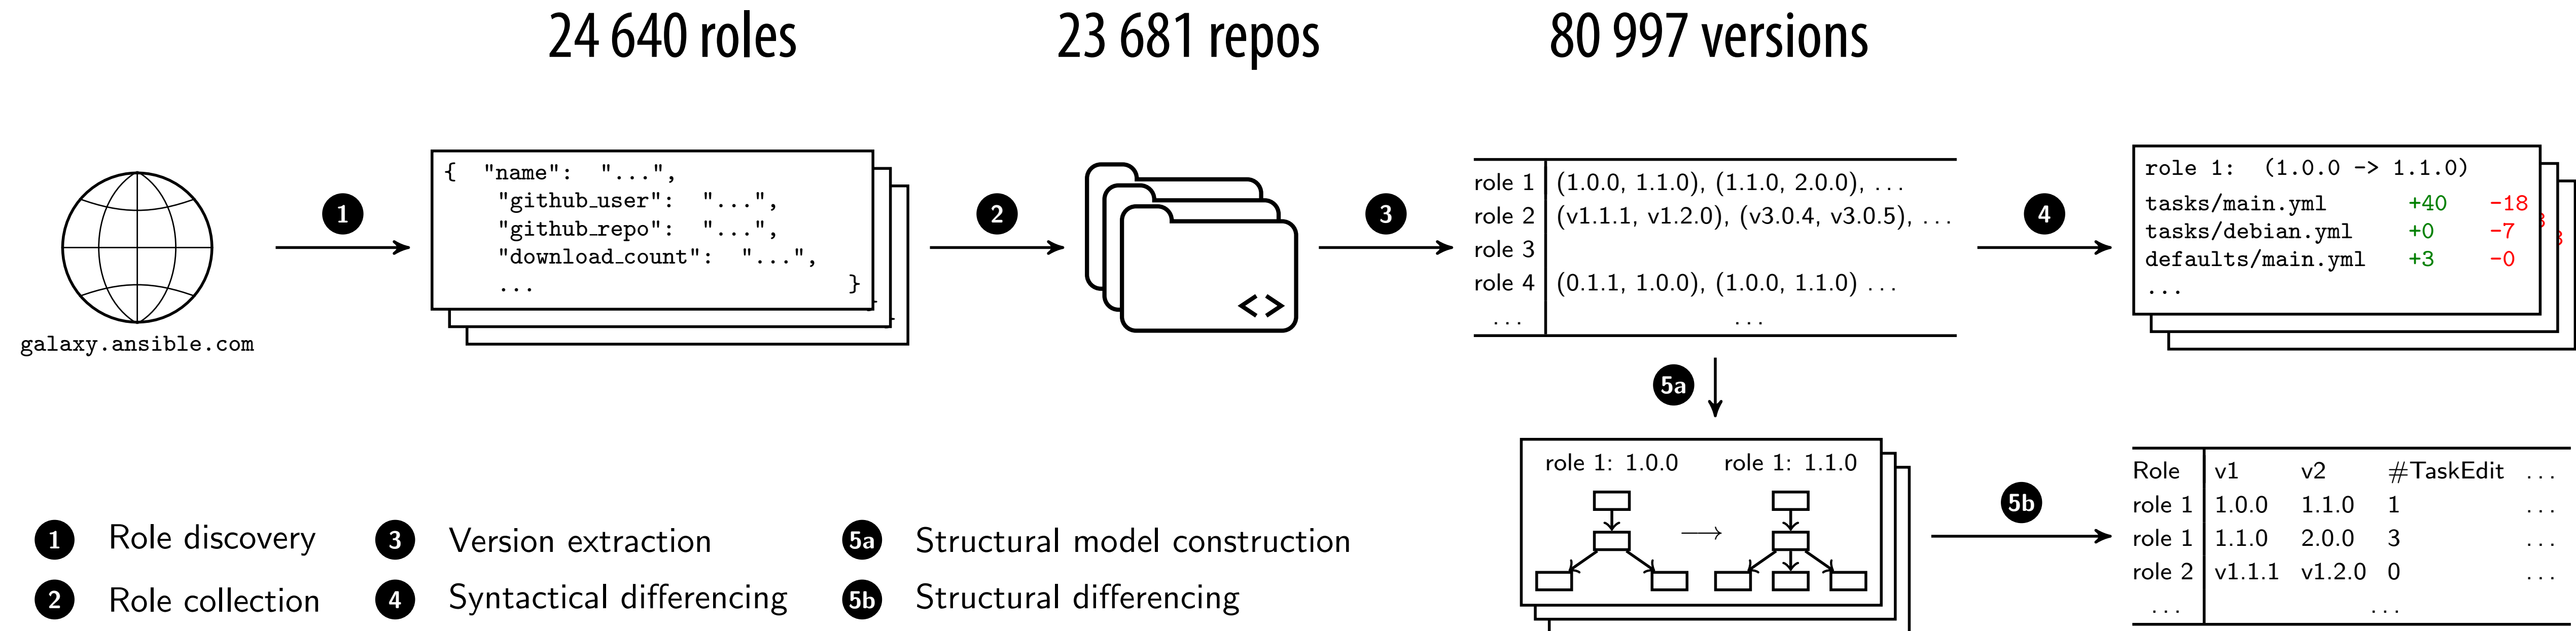
Roles 1458

 geerlingguy	nginx Nginx installation for Linux, FreeBSD and OpenBSD. balancer development load nginx proxy reverse web	build passing 4.5 / 5 Score 📄 4685056 Downloads Last Imported: 5 days ago
 nginxinc	nginx Official Ansible role for NGINX development install nginx opensource oss plus server web	build passing 3.8 / 5 Score 📄 721606 Downloads Last Imported: 2 days ago
 jdauphant	nginx Ansible role to install Nginx. web	build passing 4.8 / 5 Score 📄 257736 Downloads Last Imported: 5 months ago
 weareinteract...	nginx Installs and configures nginx nginx web	build passing 5 / 5 Score 📄 39194 Downloads Last Imported: 4 months ago
 sansible	nginx nginx installation development web	build passing 5 / 5 Score 📄 18082 Downloads Last Imported: 13 days ago
 robertdebock	nginx Install and configure nginx on your system. centos installer nginx redhat web	build passing 5 / 5 Score 📄 16758 Downloads Last Imported: 3 days ago

Ecosystem of roles
"Ansible's maven"


25k+ roles
~6k authors

Semantic Versioning in Ansible Roles



R. Opdebeeck, A. Zeraouli, C. Velázquez-Rodríguez, C. De Roover. "Does Infrastructure as Code Adhere to Semantic Versioning? An Analysis of Ansible Role Evolution", In Proc. 20th International Working Conference on Source Code Analysis and Manipulation, 2020.

Semantic Versioning in Ansible Roles



R. Opdebeeck, A. Zeraouli, C. Velázquez-Rodríguez, C. De Roover. "Does Infrastructure as Code Adhere to Semantic Versioning? An Analysis of Ansible Role Evolution", In Proc. 20th International Working Conference on Source Code Analysis and Manipulation, 2020.



Ansible Roles

Reusable Collections of Tasks

```
# roles/nginx/tasks/main.yml
---
- name: Ensure nginx is installed
  apt:
    name: nginx
    state: present
- name: Ensure nginx configuration is present
  file:
    path: "{{ item }}"
    dest: "{{ nginx_dir }}"
    loop: "{{ conf_files }}"
```

```
# roles/nginx/defaults/main.yml
---
conf_files:
  - files/nginx.conf
```

```
# roles/nginx/vars/main.yml
---
nginx_conf_directory: /etc/nginx
```

```
# roles/nginx/meta/main.yml
---
galaxy_info:
  role_name: nginx
  author: R0pdebee
  description: Installs nginx
  license: "GPL3"
  min_ansible_version: 2.4
  platforms:
    - name: Debian
      versions:
        - all
```



Ansible Roles

Reusable Collections of Tasks

```
# roles/nginx/tasks/main.yml
---
- name: Ensure nginx is installed
  apt:
    name: nginx
    state: present
- name: Ensure nginx configuration is present
  file:
    path: "{{ item }}"
    dest: "{{ nginx_dir }}"
    loop: "{{ conf_files }}"
```

```
# roles/nginx/defaults/main.yml
---
conf_files:
  - files/nginx.conf
```

```
# roles/nginx/vars/main.yml
---
nginx_conf_directory: /etc/nginx
```

```
# roles/nginx/meta/main.yml
---
galaxy_info:
  role_name: nginx
  author: R0pdebee
  description: Installs nginx
  license: "GPL3"
  min_ansible_version: 2.4
  platforms:
    - name: Debian
      versions:
        - all
```

Tasks
Basic building blocks



Ansible Roles

Reusable Collections of Tasks

```
# roles/nginx/tasks/main.yml
---
- name: Ensure nginx is installed
  apt:
    name: nginx
    state: present
- name: Ensure nginx configuration is present
  file:
    path: "{{ item }}"
    dest: "{{ nginx_dir }}"
    loop: "{{ conf_files }}"
```

```
# roles/nginx/defaults/main.yml
---
conf_files:
  - files/nginx.conf
```

```
# roles/nginx/vars/main.yml
---
nginx_conf_directory: /etc/nginx
```

```
# roles/nginx/meta/main.yml
---
galaxy_info:
  role_name: nginx
  author: R0pdebee
  description: Installs nginx
  license: "GPL3"
  min_ansible_version: 2.4
  platforms:
    - name: Debian
      versions:
        - all
```

Variables



Ansible Roles

Reusable Collections of Tasks

```
# roles/nginx/tasks/main.yml
---
- name: Ensure nginx is installed
  apt:
    name: nginx
    state: present
- name: Ensure nginx configuration is present
  file:
    path: "{{ item }}"
    dest: "{{ nginx_dir }}"
    loop: "{{ conf_files }}"
```

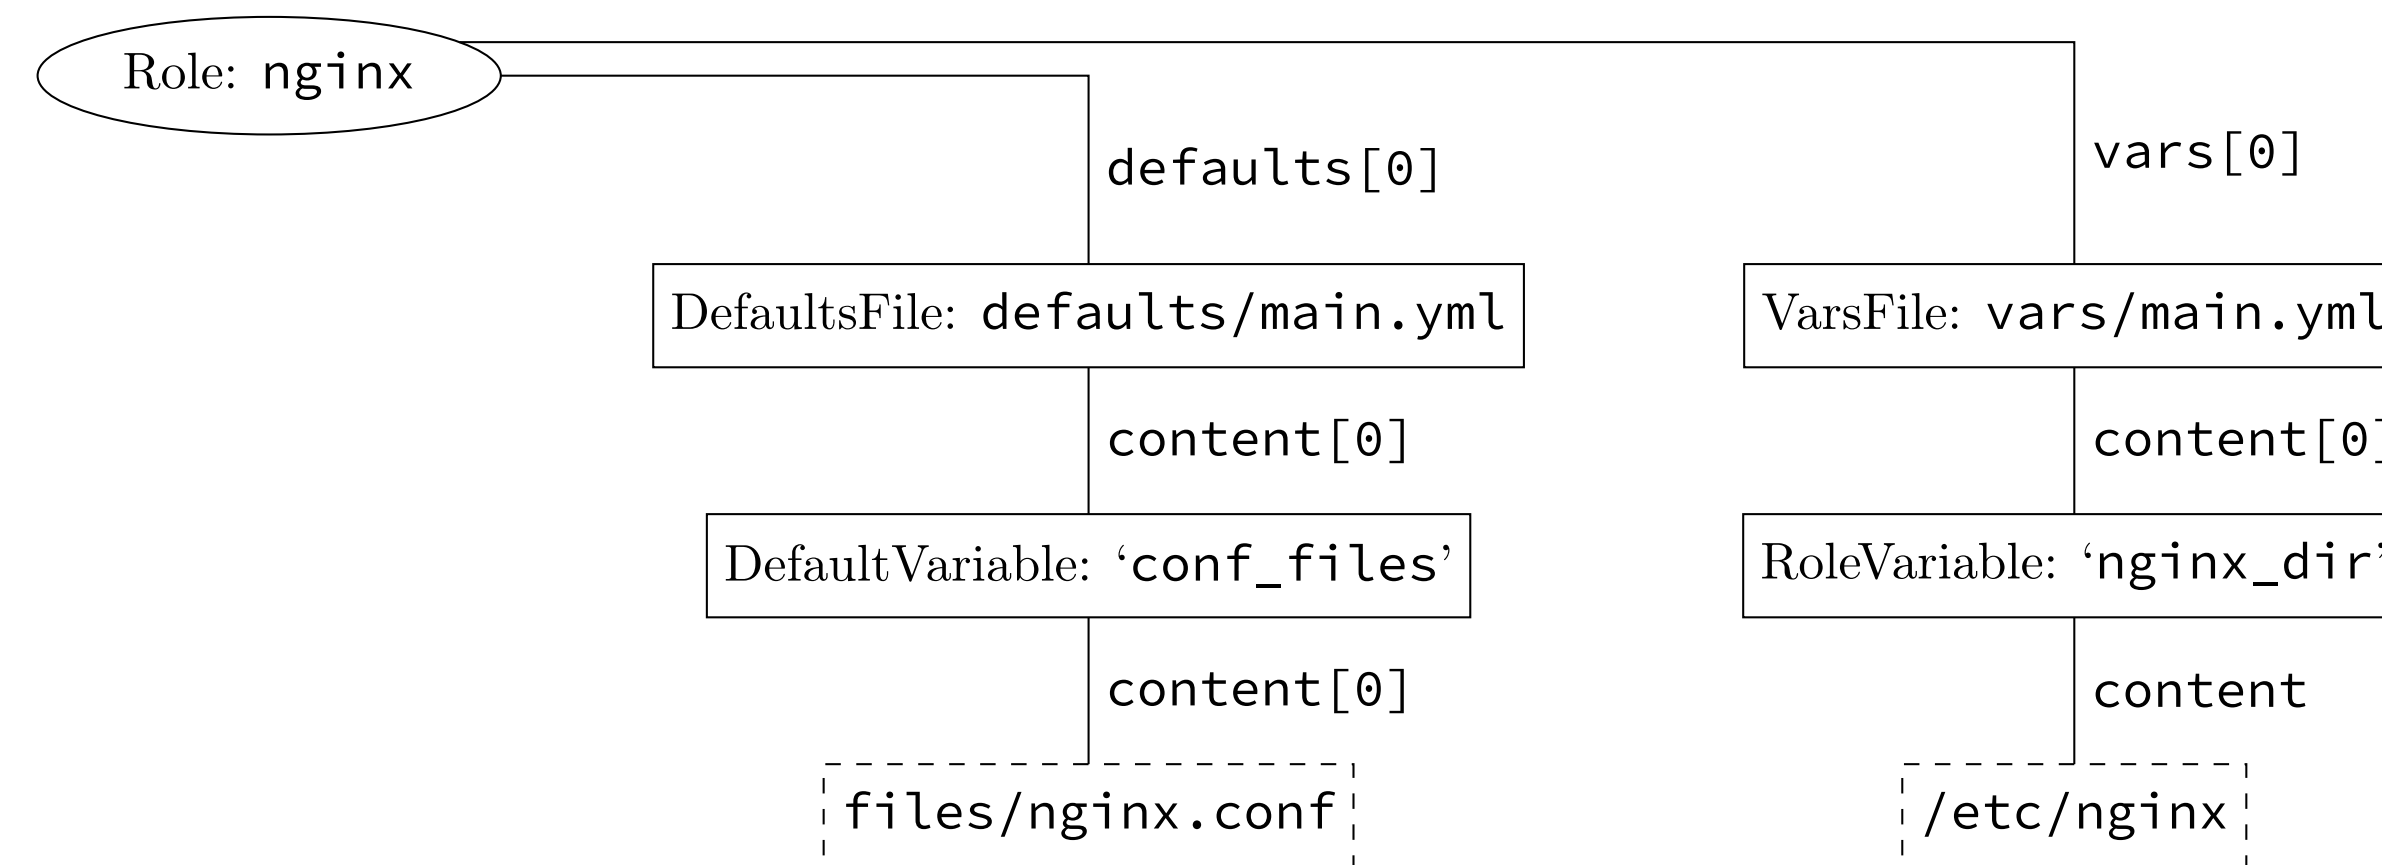
```
# roles/nginx/defaults/main.yml
---
conf_files:
  - files/nginx.conf
```

```
# roles/nginx/vars/main.yml
---
nginx_conf_directory: /etc/nginx
```

```
# roles/nginx/meta/main.yml
---
galaxy_info:
  role_name: nginx
  author: R0pdebee
  description: Installs nginx
  license: "GPL3"
  min_ansible_version: 2.4
  platforms:
    - name: Debian
      versions:
        - all
```

**Role metadata
Information for Galaxy**

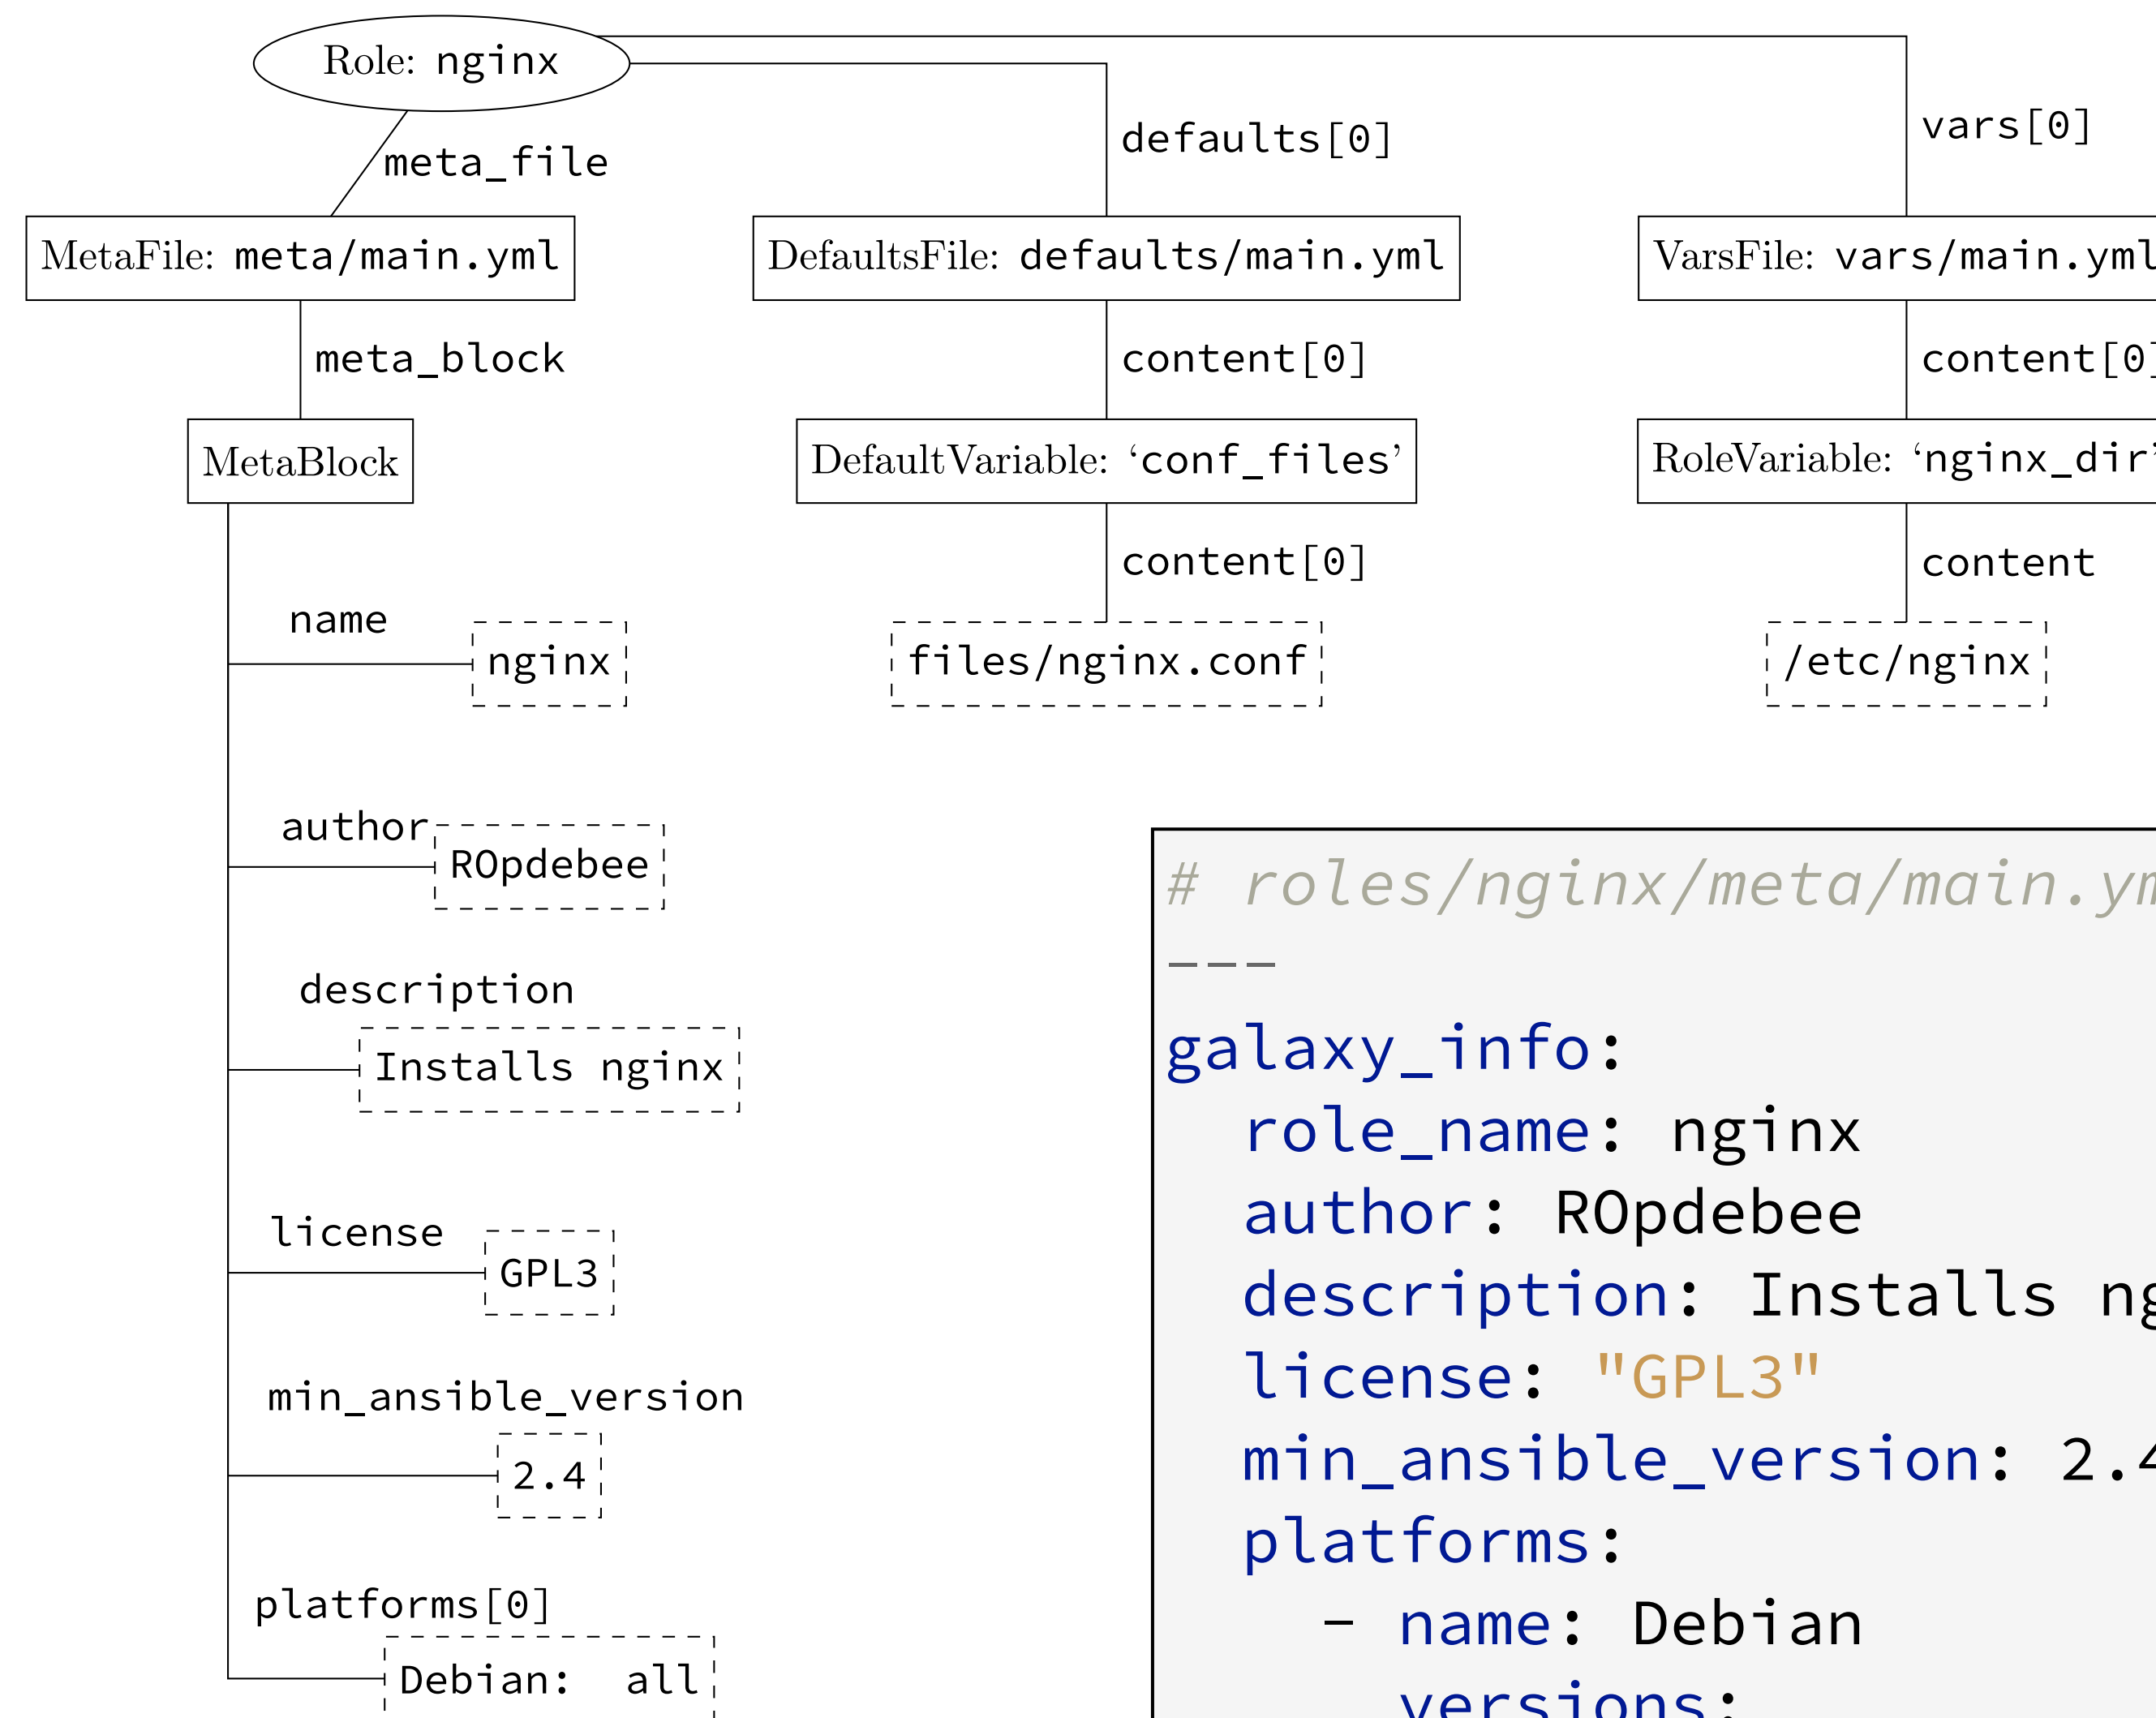
Structural Representation of Ansible Roles



```
# roles/nginx/defaults/main.yml  
---  
conf_files:  
  - files/nginx.conf
```

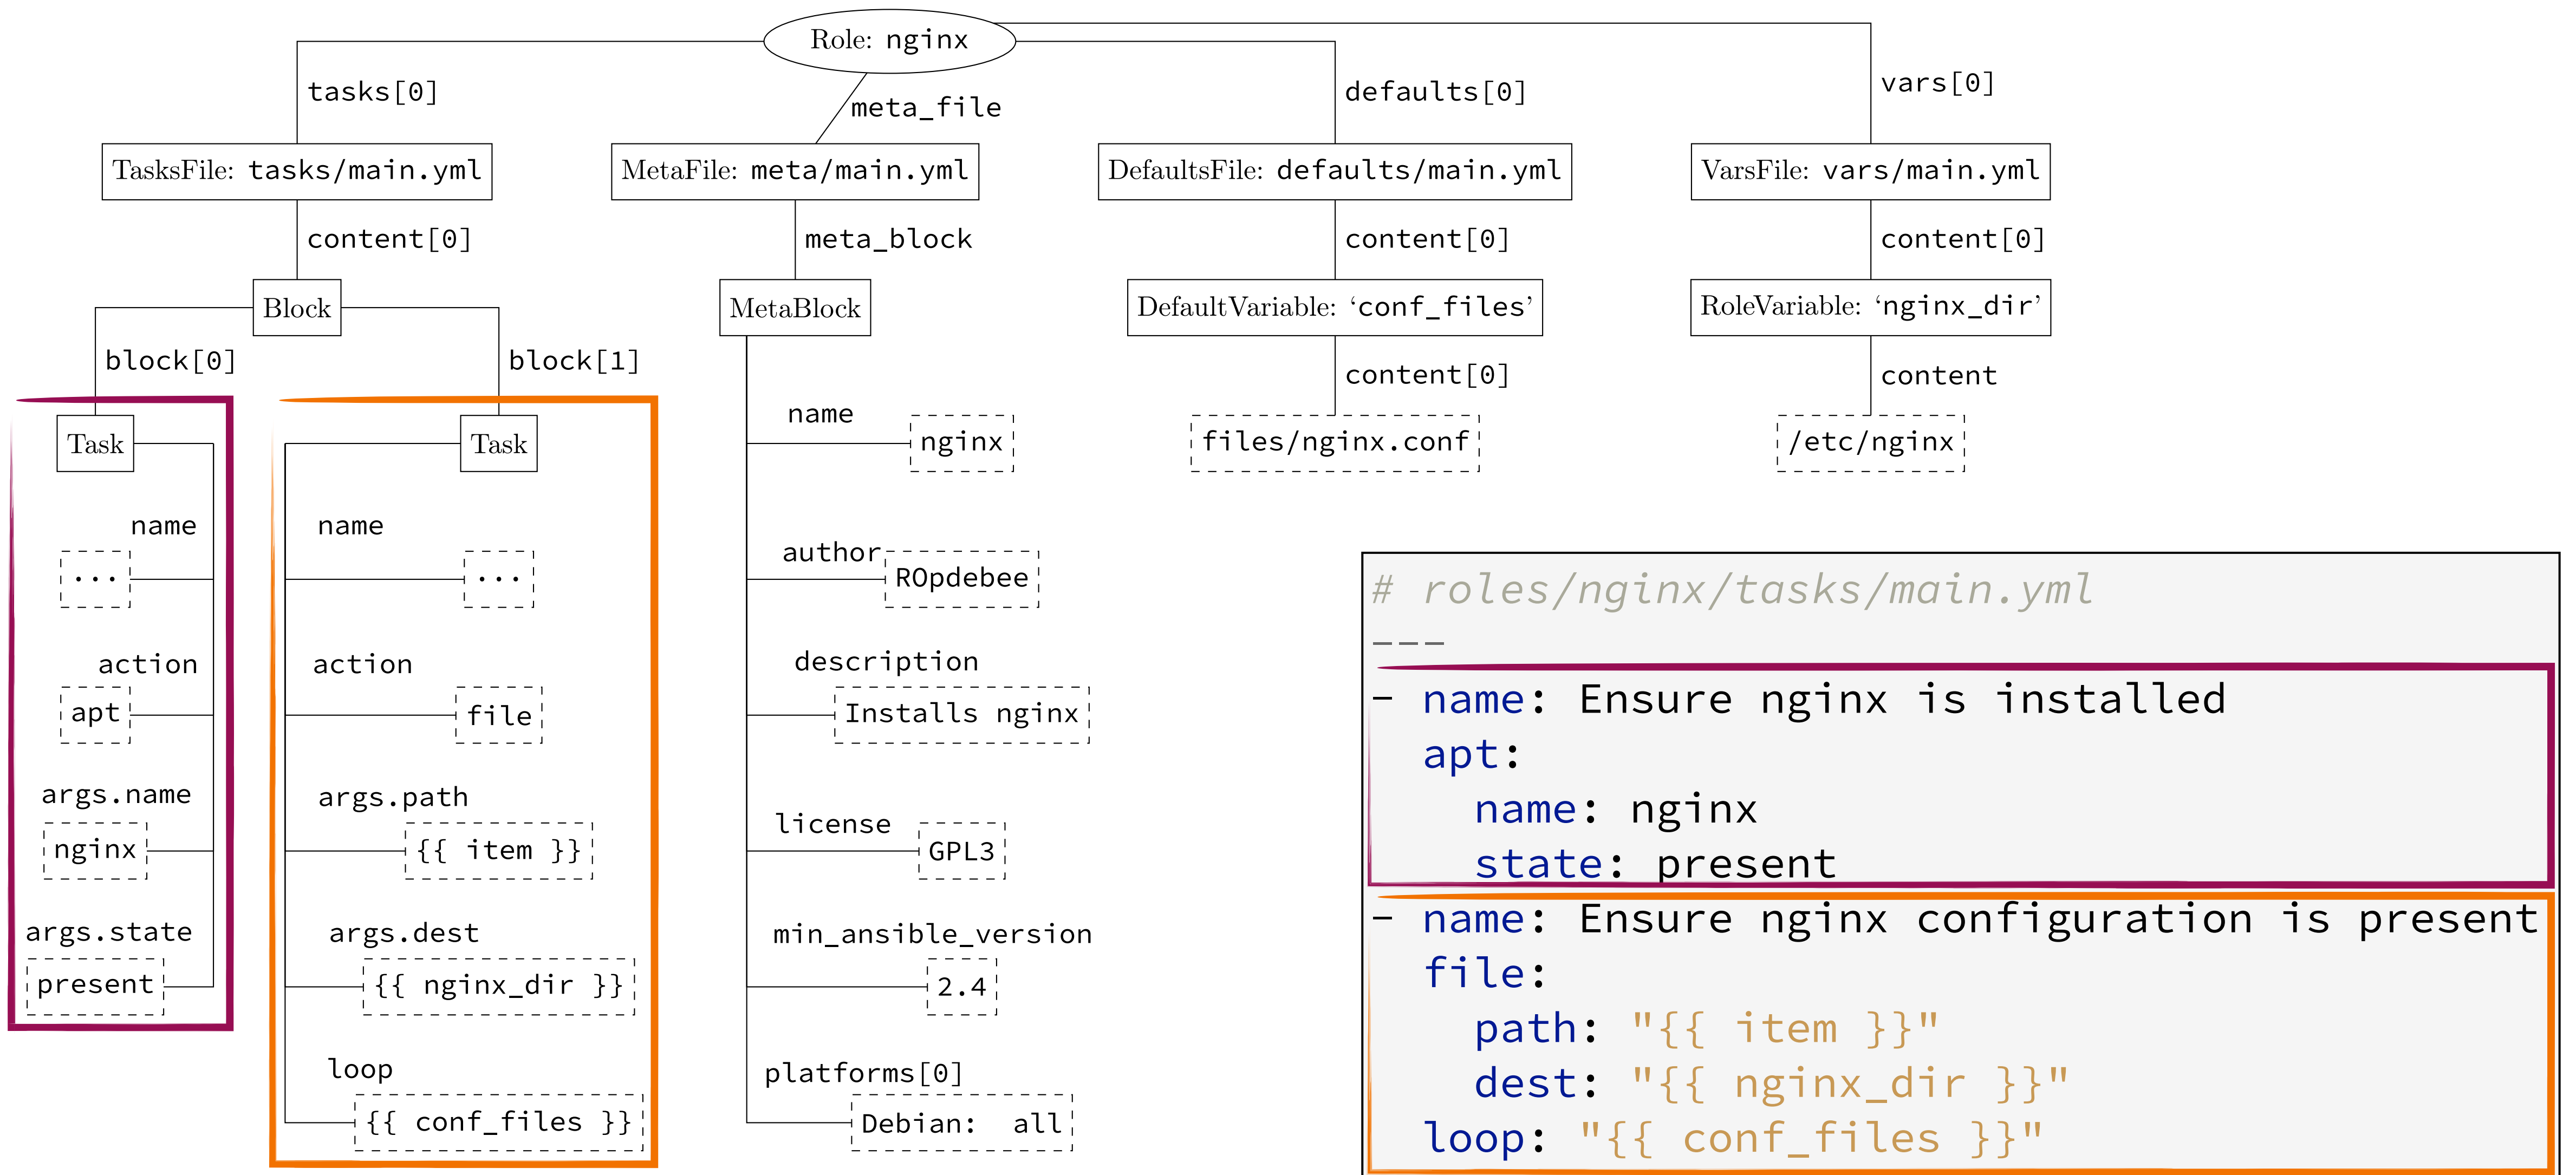
```
# roles/nginx/vars/main.yml  
---  
nginx_conf_directory: /etc/nginx
```

Structural Representation of Ansible Roles



```
# roles/nginx/meta/main.yml
---
galaxy_info:
  role_name: nginx
  author: R0pdebee
  description: Installs nginx
  license: "GPL3"
  min_ansible_version: 2.4
  platforms:
    - name: Debian
      versions:
        - all
```

Structural Representation of Ansible Roles



Change Distilling

Extracting Changes After the Fact

Empirical Software Engineering (2019) 24: 491–535
https://doi.org/10.1007/s10664-018-9644-3



Querying distilled code changes to extract executable transformations

Reinout Stevens¹ · Tim Molderez² · Coen De Roover²

Published online: 30 August 2018
© Springer Science+Business Media, LLC, part of Springer Nature 2018

Abstract

Change distilling algorithms compute a sequence of fine-grained changes that, when executed in order, transform a given source AST into a given target AST. The resulting change sequences are used in the field of mining software repositories to study source code evolution. Unfortunately, detecting and specifying source code evolutions in such a change sequence is cumbersome. We therefore introduce a tool-supported approach that identifies minimal executable subsequences in a sequence of distilled changes that implement a particular evolution pattern, specified in terms of intermediate states of the AST that undergoes each change. This enables users to describe the effect of multiple changes, irrespective of their execution order, while ensuring that different change sequences that implement the same code evolution are recalled. Correspondingly, our evaluation is two-fold. We show that our approach is able to recall different implementation variants of the same source code evolution in histories of different software projects. We also evaluate the expressiveness and ease-of-use of our approach in a user study.

Keywords Change distilling · Change querying · Logic meta-programming

1 Introduction

The use of a Version Control System (VCS) has become an industry best practice for developing software. Researchers in the field of mining software repositories (MSR) leverage the resulting revision histories to study the evolution of software systems. However, most VCSs

Communicated by: Gabriele Bavota and Andrian Marcus

✉ Tim Molderez
tim.molderez@vub.be

Reinout Stevens
reinout@reinoutstevens.be

Coen De Roover
coen.de.roover@vub.be

¹ Maxflow BVBA, Leuven, Belgium

² Software Languages Lab, Vrije Universiteit Brussel, Ixelles, Belgium



GumTree

ChangeDistiller

ChangeNodes

Fine-grained and Accurate Source Code Differencing

Jean-Rémy Falleri
Univ. Bordeaux,
LaBRI, UMR 5800
F-33400, Talence, France
falleri@labri.fr

Floréal Morandat
Univ. Bordeaux,
LaBRI, UMR 5800
F-33400, Talence, France
fmoranda@labri.fr

Xavier Blanc
Univ. Bordeaux,
LaBRI, UMR 5800
F-33400, Talence, France
xblanc@labri.fr

Matias Martinez
INRIA and University of Lille,
France
matias.martinez@inria.fr

Martin Monperrus
INRIA and University of Lille,
France
martin.monperrus@inria.fr

ABSTRACT

At the heart of software evolution is a sequence of edit actions, called an *edit script*, made to a source code file. Since software systems are stored version by version, the edit script has to be computed from these versions, which is known as a complex task. Existing approaches usually compute edit scripts at the text granularity with only *add line* and *delete line* actions. However, inferring syntactic changes from such an edit script is hard. Since moving code is a frequent action performed when editing code and it should also be taken into account. In this paper, we tackle these issues by introducing an algorithm computing edit scripts at the abstract syntax tree granularity including move actions. Our objective is to compute edit scripts that are short and close to the original developer intent. Our algorithm is implemented in a freely-available and extensible tool that has been intensively validated.

Categories and Subject Descriptors: D.2.3 [Software Engineering]: Coding Tools and Techniques

General Terms: Algorithms, Experimentation

Keywords: Software evolution, Program comprehension, Tree differencing, AST.

1. INTRODUCTION

The first law of software evolution states that almost all software systems have to evolve to be satisfactory [19]. Since this law was formulated, many studies have been performed to better understand how software systems evolve, and forms what is called the *software evolution* research field [21].

There is global software evolution (e.g. evolution of requirements, of execution environments, ...) and local software evolution (evolution of source code files). In this paper, we focus on the latter, that is on understanding how source code files evolve. In particular, we focus on *edit scripts*, that are sequences of edit actions made to a source code file. Usually, since software is stored in version control systems, edit scripts Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. ASE '14, September 15–19, 2014, Västerås, Sweden. Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM 978-1-4503-3013-8/14/09...\$15.00. http://dx.doi.org/10.1145/2642937.2642982.

are computed between two versions of a same file. The goal of an edit script is to accurately reflect the actual change that has been performed on a file.

Edit scripts are used by developers on a daily basis. For example, the Unix *diff* tool takes as input two versions of a source code file and performs the Myers algorithm [24] at the text line granularity and returns an edit script indicating which lines have been added or deleted. However, the limitations of *diff* are twofold. First, it only computes additions and deletions and does not consider other kinds of edit actions such as update and move. Second, it works at a granularity (the text line) that is both coarse grain and not aligned with the source code structure: the abstract syntax tree.

To overcome this main limitation, there are algorithms that can work at the abstract syntax tree (AST) level [13]. The main advantage in using the AST granularity is that the edit script directly refers to the structure of the code. For instance, if an edit action is the addition of a new function node, it clearly means that a new function has been added in the code. Despite several key contributions (e.g. [13]), the problem of computing AST edit scripts is still open, with two main challenges: handling move actions, and scaling to fine-grained ASTs with thousands of nodes¹. This is where this paper makes a contribution.

To design our novel algorithm, we take the viewpoint of the developer: she is never interested in the theoretical shortest edit script. She is rather interested in having an edit script that reflects well the actual changes that happened. Thus our objective is not to find the shortest sequence of actions between two versions, but a sequence that reflects well the developer intent. Consequently, we devise an algorithm based on heuristics that contain pragmatic rules on what a good edit script is, and as importantly, that is efficient and scales to large ASTs. This algorithm has been implemented within a freely-available and extensible tool².

To sum up, our contributions are:

- a novel efficient AST differencing algorithm that takes into account move actions, and its implementation;

¹The best known algorithm with add, delete and update actions has a $O(n^3)$ time complexity with n being the number of nodes of the AST [27]. Computing the minimum edit script that can include move node actions is known to be NP-hard [4].

²github.com/jrfaller/gumtree

IEEE TRANSACTIONS ON SOFTWARE ENGINEERING, VOL. 33, NO. 11, NOVEMBER 2007

725

Change Distilling: Tree Differencing for Fine-Grained Source Code Change Extraction

Beat Fluri, *Student Member, IEEE*, Michael Würsch, *Student Member, IEEE*,
Martin Pinzger, *Member, IEEE*, and Harald C. Gall, *Member, IEEE*

Abstract—A key issue in software evolution analysis is the identification of particular changes that occur across several versions of a program. We present *change distilling*, a tree differencing algorithm for fine-grained source code change extraction. For that, we have improved the existing algorithm by Chawathe et al. for extracting changes in hierarchically structured data [8]. Our algorithm extracts changes by finding both a match between the nodes of the compared two abstract syntax trees and a minimum edit script that can transform one tree into the other given the computed matching. As a result, we can identify fine-grained change types between program versions according to our *taxonomy of source code changes*. We evaluated our change distilling algorithm with a benchmark that we developed, which consists of 1,064 manually classified changes in 219 revisions of eight methods from three different open source projects. We achieved significant improvements in extracting types of source code changes: Our algorithm approximates the minimum edit script 45 percent better than the original change extraction approach by Chawathe et al. We are able to find all occurring changes and almost reach the minimum conforming edit script, that is, we reach a mean absolute percentage error of 34 percent, compared to the 79 percent reached by the original algorithm. The paper describes both our change distilling algorithm and the results of our evaluation.

Index Terms—Source code change extraction, tree-differencing algorithms, software repositories, software evolution analysis.

1 INTRODUCTION

SINCE Lehman's Laws of Program Evolution from the 1980s [25], it has been well understood that software has to be adapted to changing requirements and environments or it becomes progressively less useful. Change is broadly accepted as a crucial part of a software's life cycle. As a consequence, in recent years, several techniques and tools have been developed to aid software engineers in maintaining and evolving large complex software systems. For instance, Ying et al. or Zimmermann et al. developed approaches that guide programmers along related changes by telling them "programmers who changed these functions also changed..." [45], [47]. The Hipikat tool of Cubranić et al. used project history information to provide recommendations for a modification task [9]. Gall et al. detected possible maintainability hot spots by analyzing cochange relationships of modules [13].

We argue that such techniques and tools are valuable but suffer from the low quality of information available for changes. Typically, such information, in particular for source code, is stored by versioning systems (for example, CVS or Subversion). They keep track of changes by storing the text lines *added* and/or *deleted* from a particular file. Structural changes in the source code are not considered at all.

• The authors are with the Department of Informatics, University of Zurich, Binzmühlstrasse 14, CH-8050 Zurich, Switzerland.
E-mail: {fluri, wuersch, pinzger, gall}@ifi.uzh.ch.

Manuscript received 15 Jan. 2007; revised 13 July 2007; accepted 23 July 2007; published online 3 Aug. 2007.

Recommended for acceptance by H. Muller.
For information on obtaining reprints of this article, please send e-mail to: ts@computer.org, and reference IEEECS Log Number TSE-0012-0107.

NP-hard [4].

²github.com/jrfaller/gumtree

0098-5589/07/\$25.00 © 2007 IEEE. Published by the IEEE Computer Society.
Authorized licensed use limited to: Vrije Universiteit Brussel. Downloaded on November 30, 2020 at 16:05:15 UTC from IEEE Xplore. Restrictions apply.

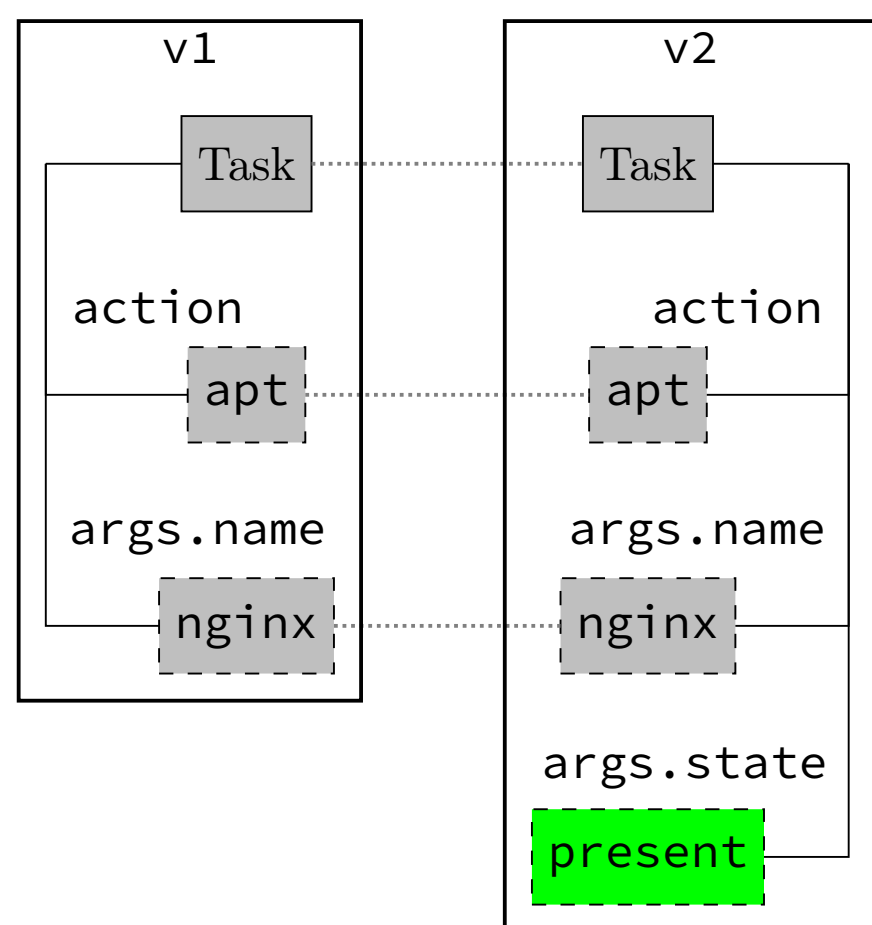
More sophisticated approaches are able to narrow down changes to the method level, but fail in further qualifying changes such as the addition of a method invocation in the else branch of an if-statement. Furthermore, a classification of changes according to their impact on other source code entities is missing. In particular, the latter information is important to improving the quality of software evolution results and, as a consequence, to providing better support for programmers and system analysts.

Since source code can be represented as abstract syntax trees (ASTs), tree differencing can be used to extract detailed change information. This approach is promising because exact information on each entity and statement is available in an AST. In our previous work [12], we built a *taxonomy of source code changes* that defines source code changes according to tree edit operations in the AST and classifies each change type with a *significance level*. The level expresses how strongly a change may impact other source code entities and whether a change may be functionality modifying or functionality preserving. In our taxonomy, we focus on object-oriented programming languages (OOPLs) and Java in particular. By adjusting the change type extraction, the taxonomy can also be used for other OOPLs. In total, our taxonomy defines 35 change types.

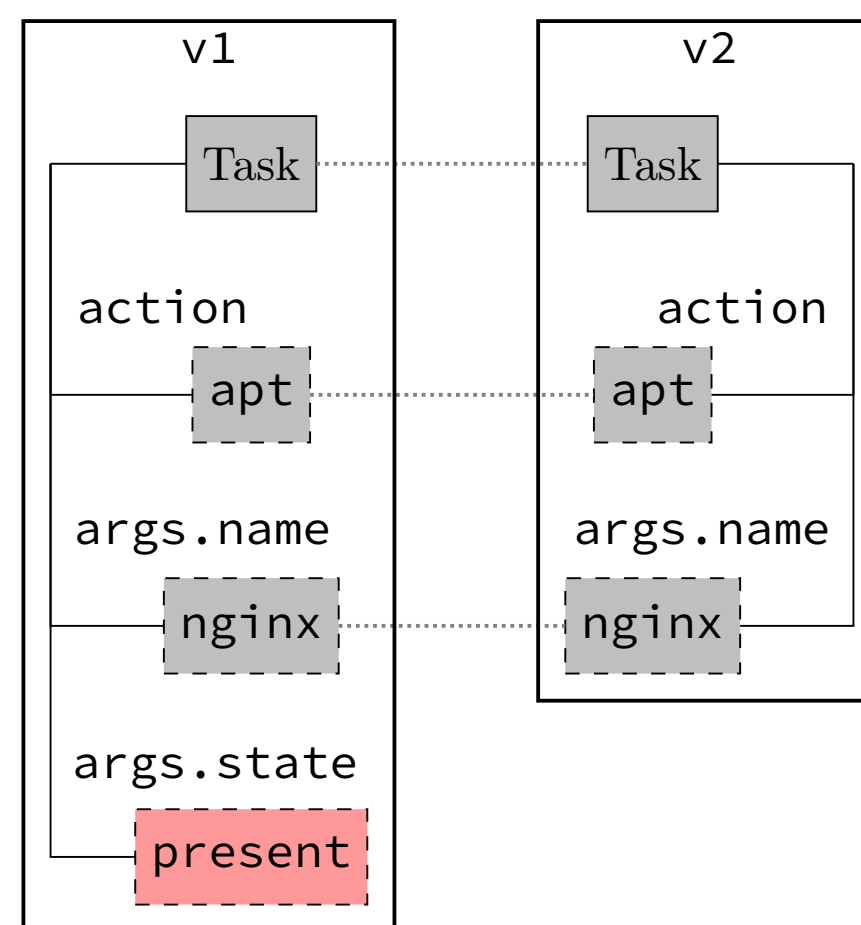
In this paper, we present *change distilling*, a tree-differencing algorithm for fine-grained source code change extraction. For that, we improved the existing algorithm for extracting changes in hierarchically structured data by Chawathe et al. [8]. This algorithm finds changes according to basic tree edit operations such as *insert*, *delete*, *move*, or *update* of tree nodes.

Change Distilling of Ansible Roles

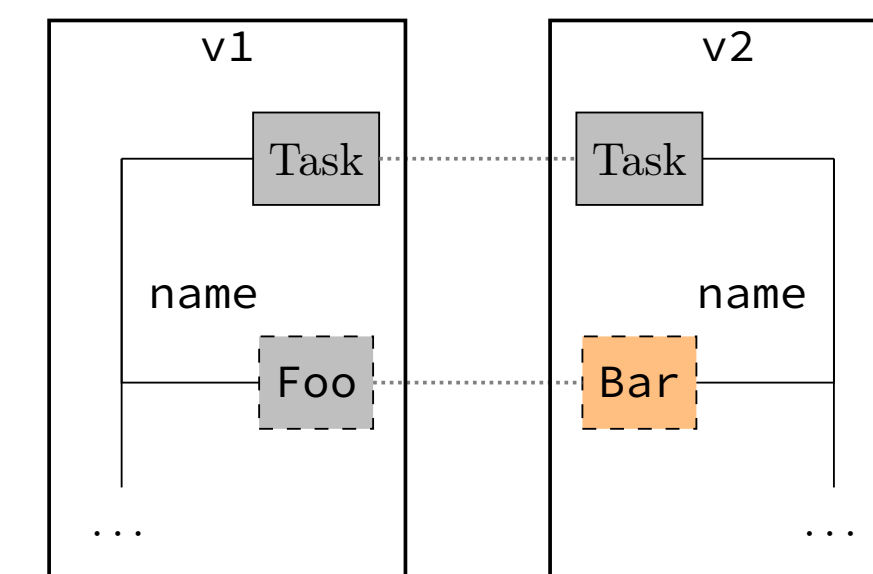
Additions



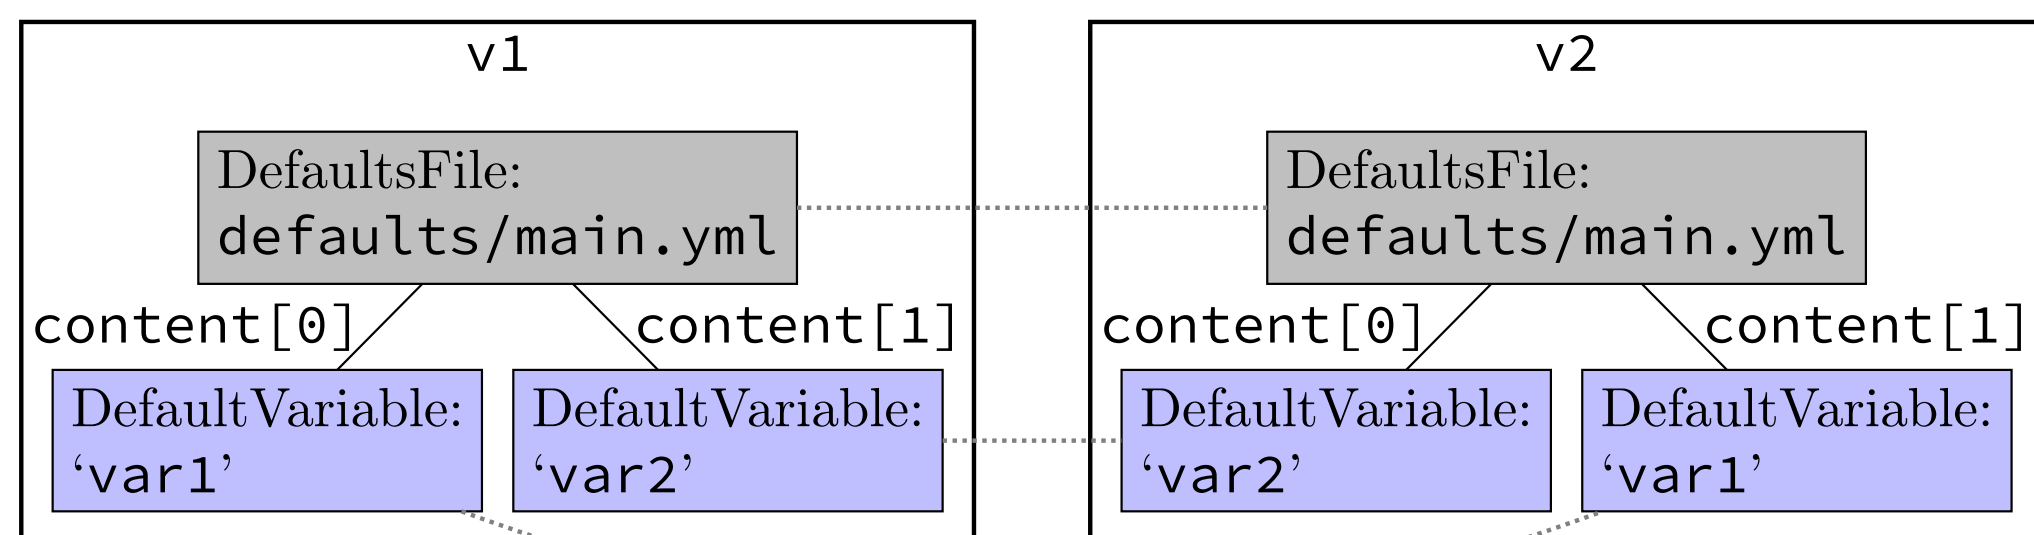
Removals



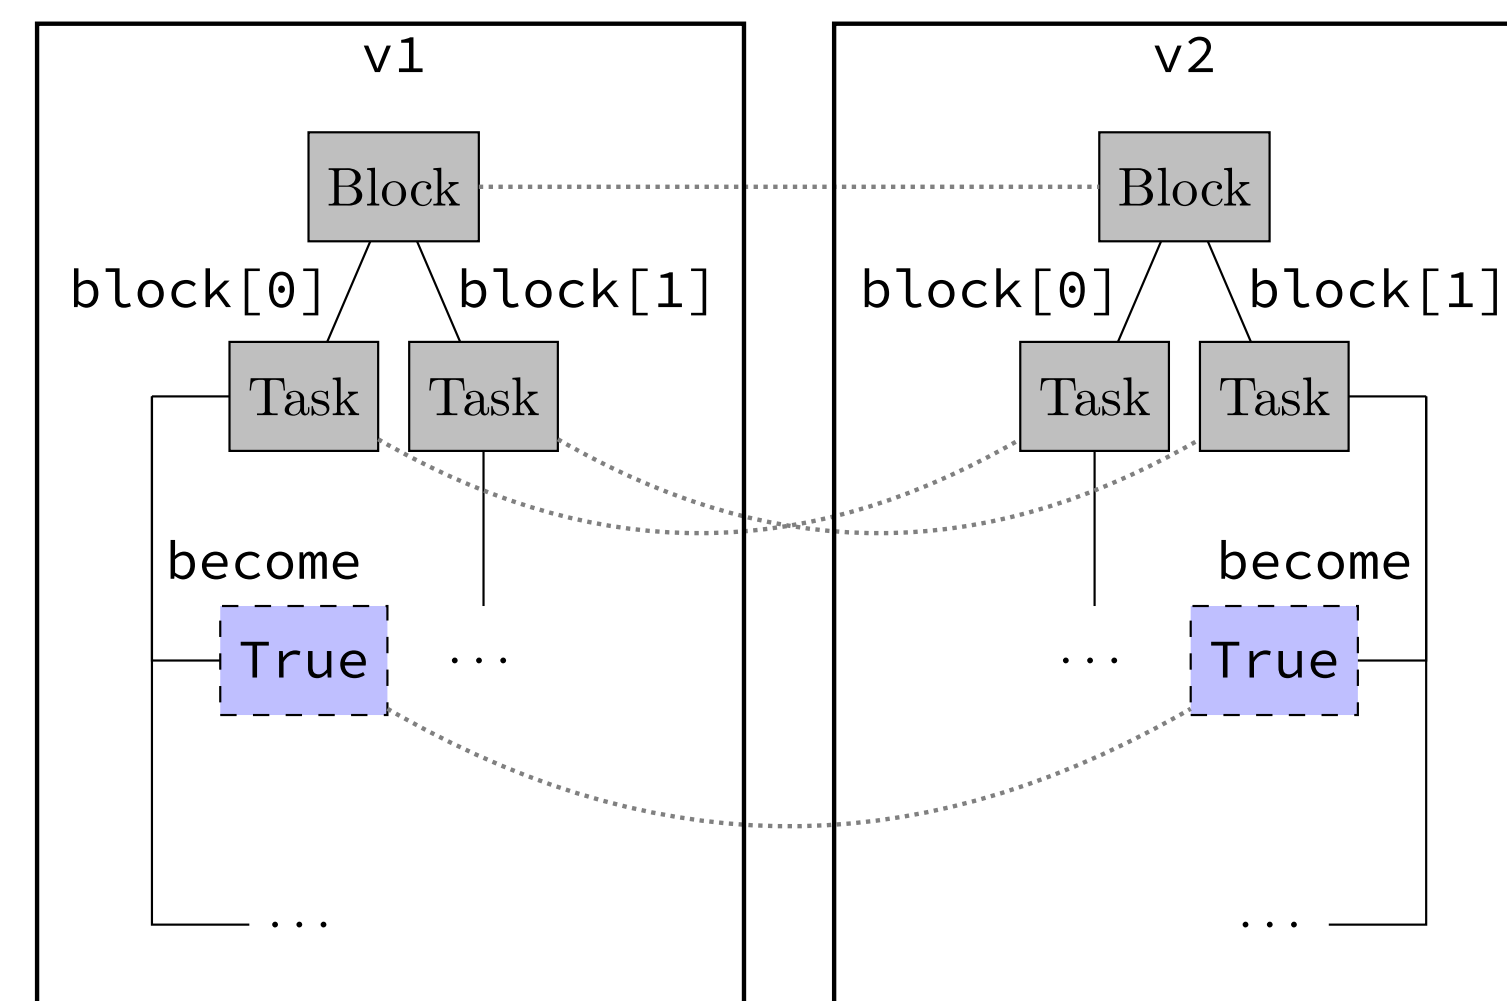
Updates



Local relocation

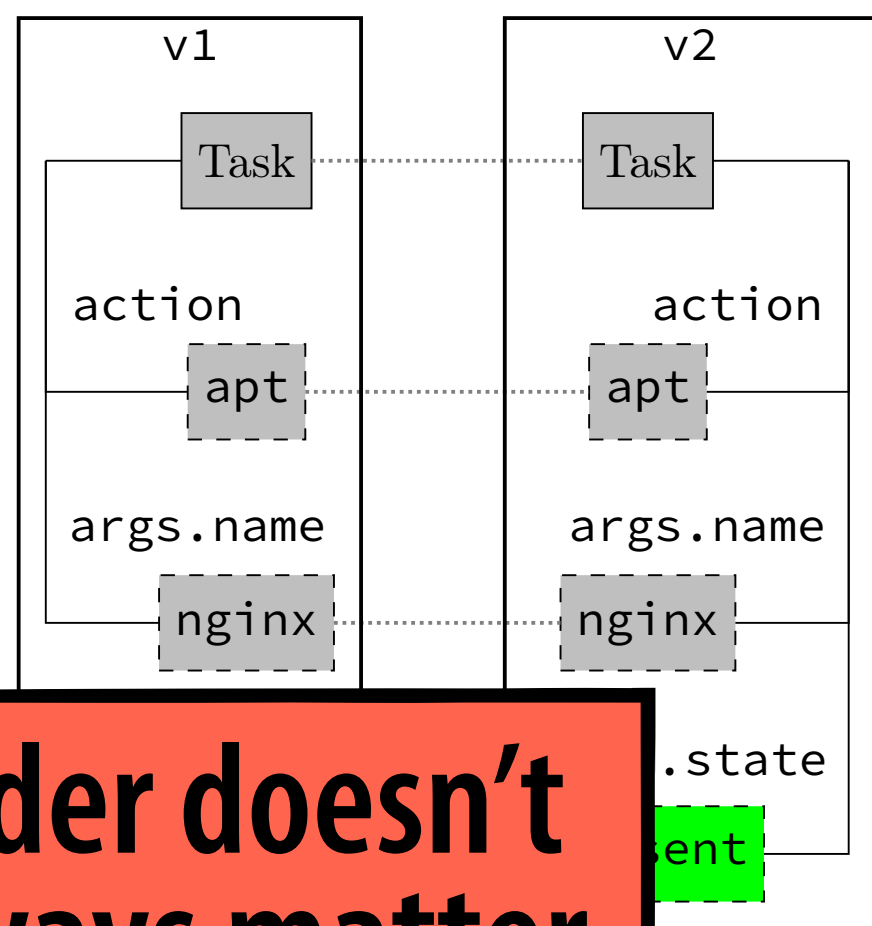


Global relocation



Change Distilling of Ansible Roles

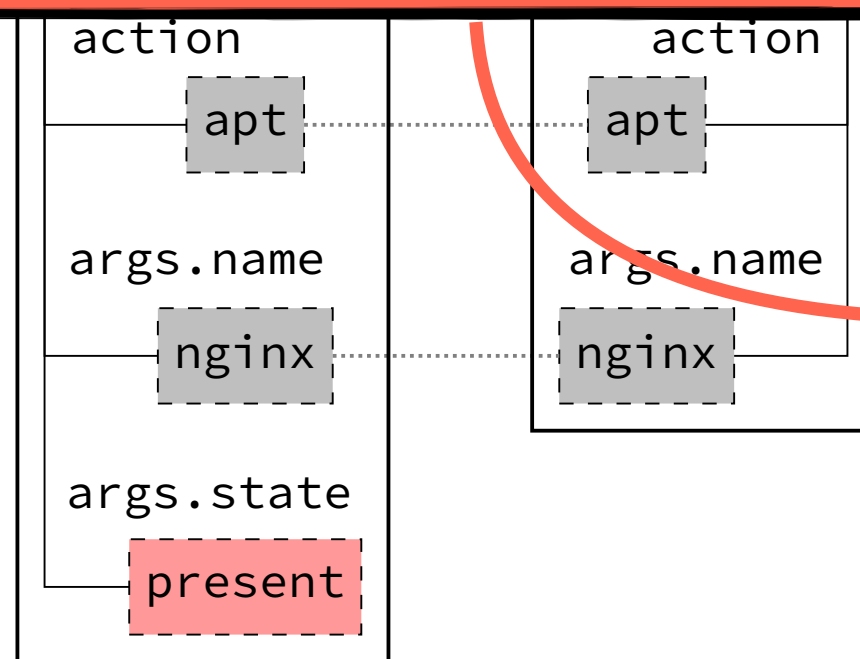
Additions



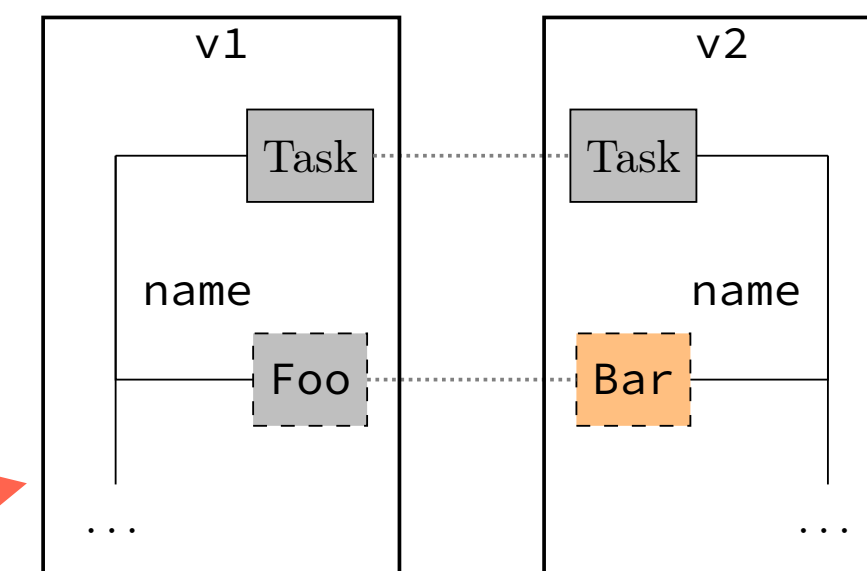
Order doesn't always matter

Removals

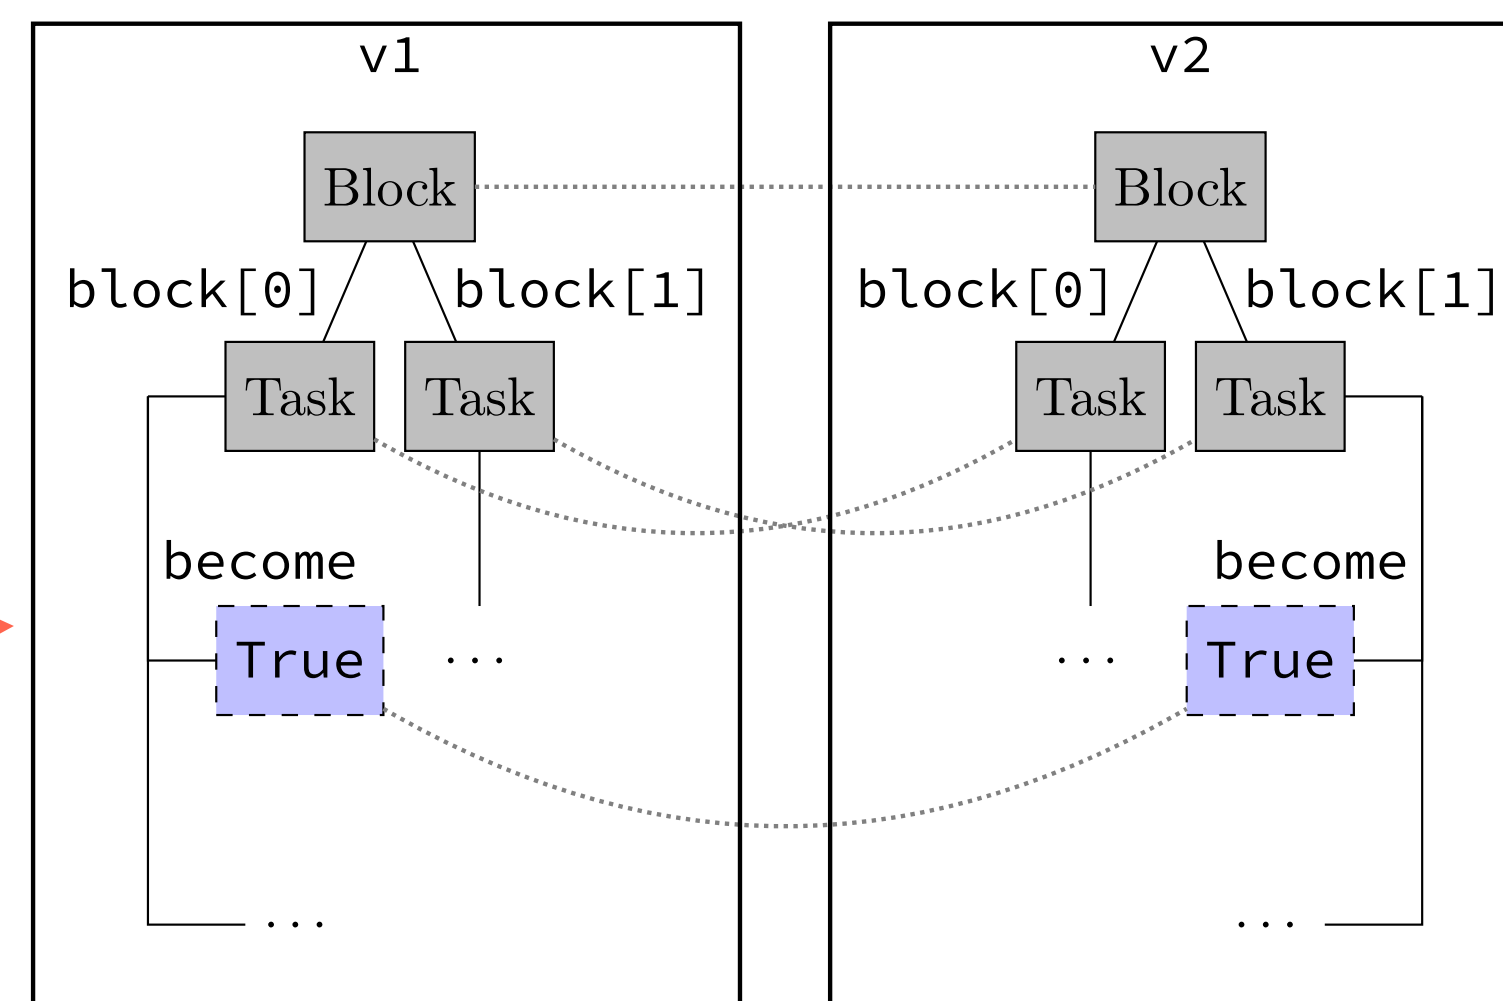
Name has no semantic relevance
Update shouldn't matter



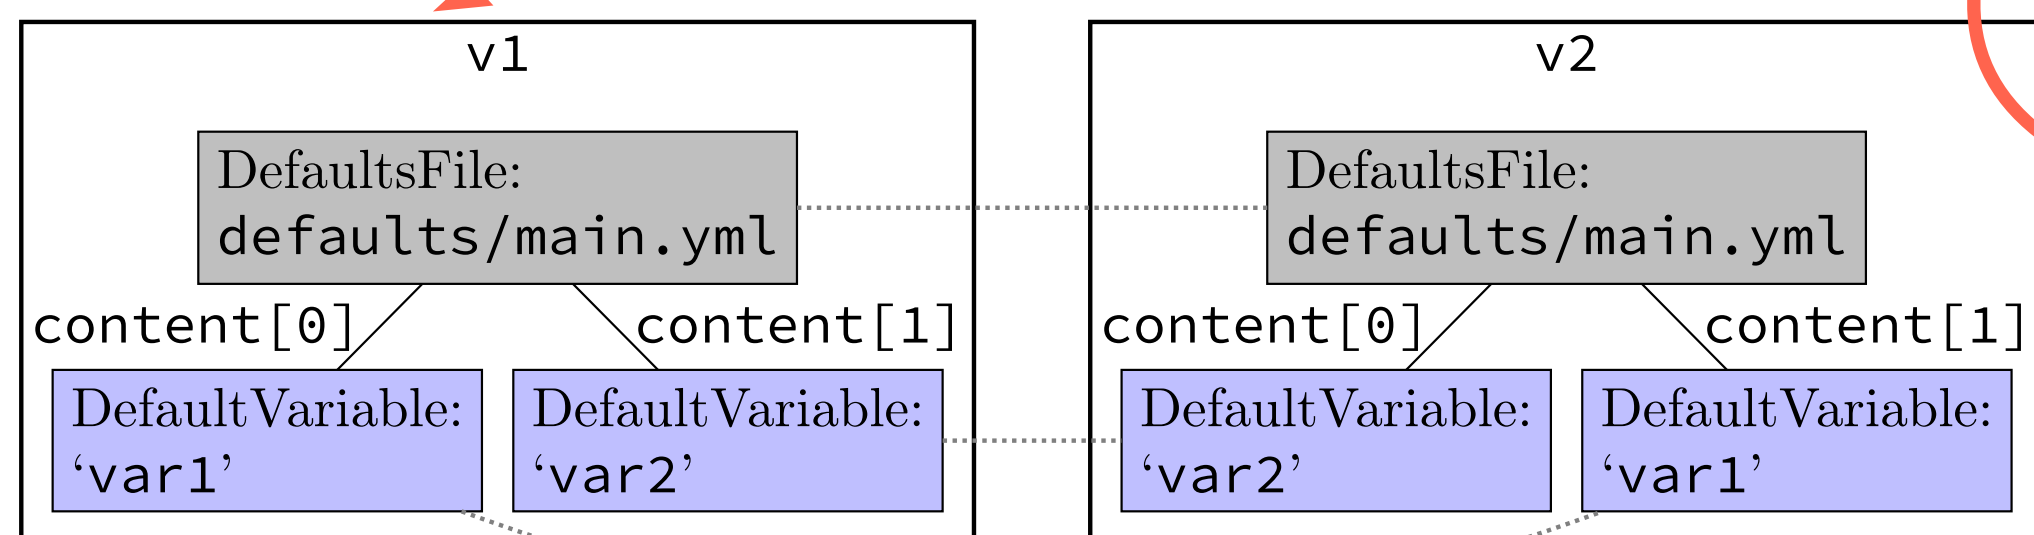
Updates



Global relocation

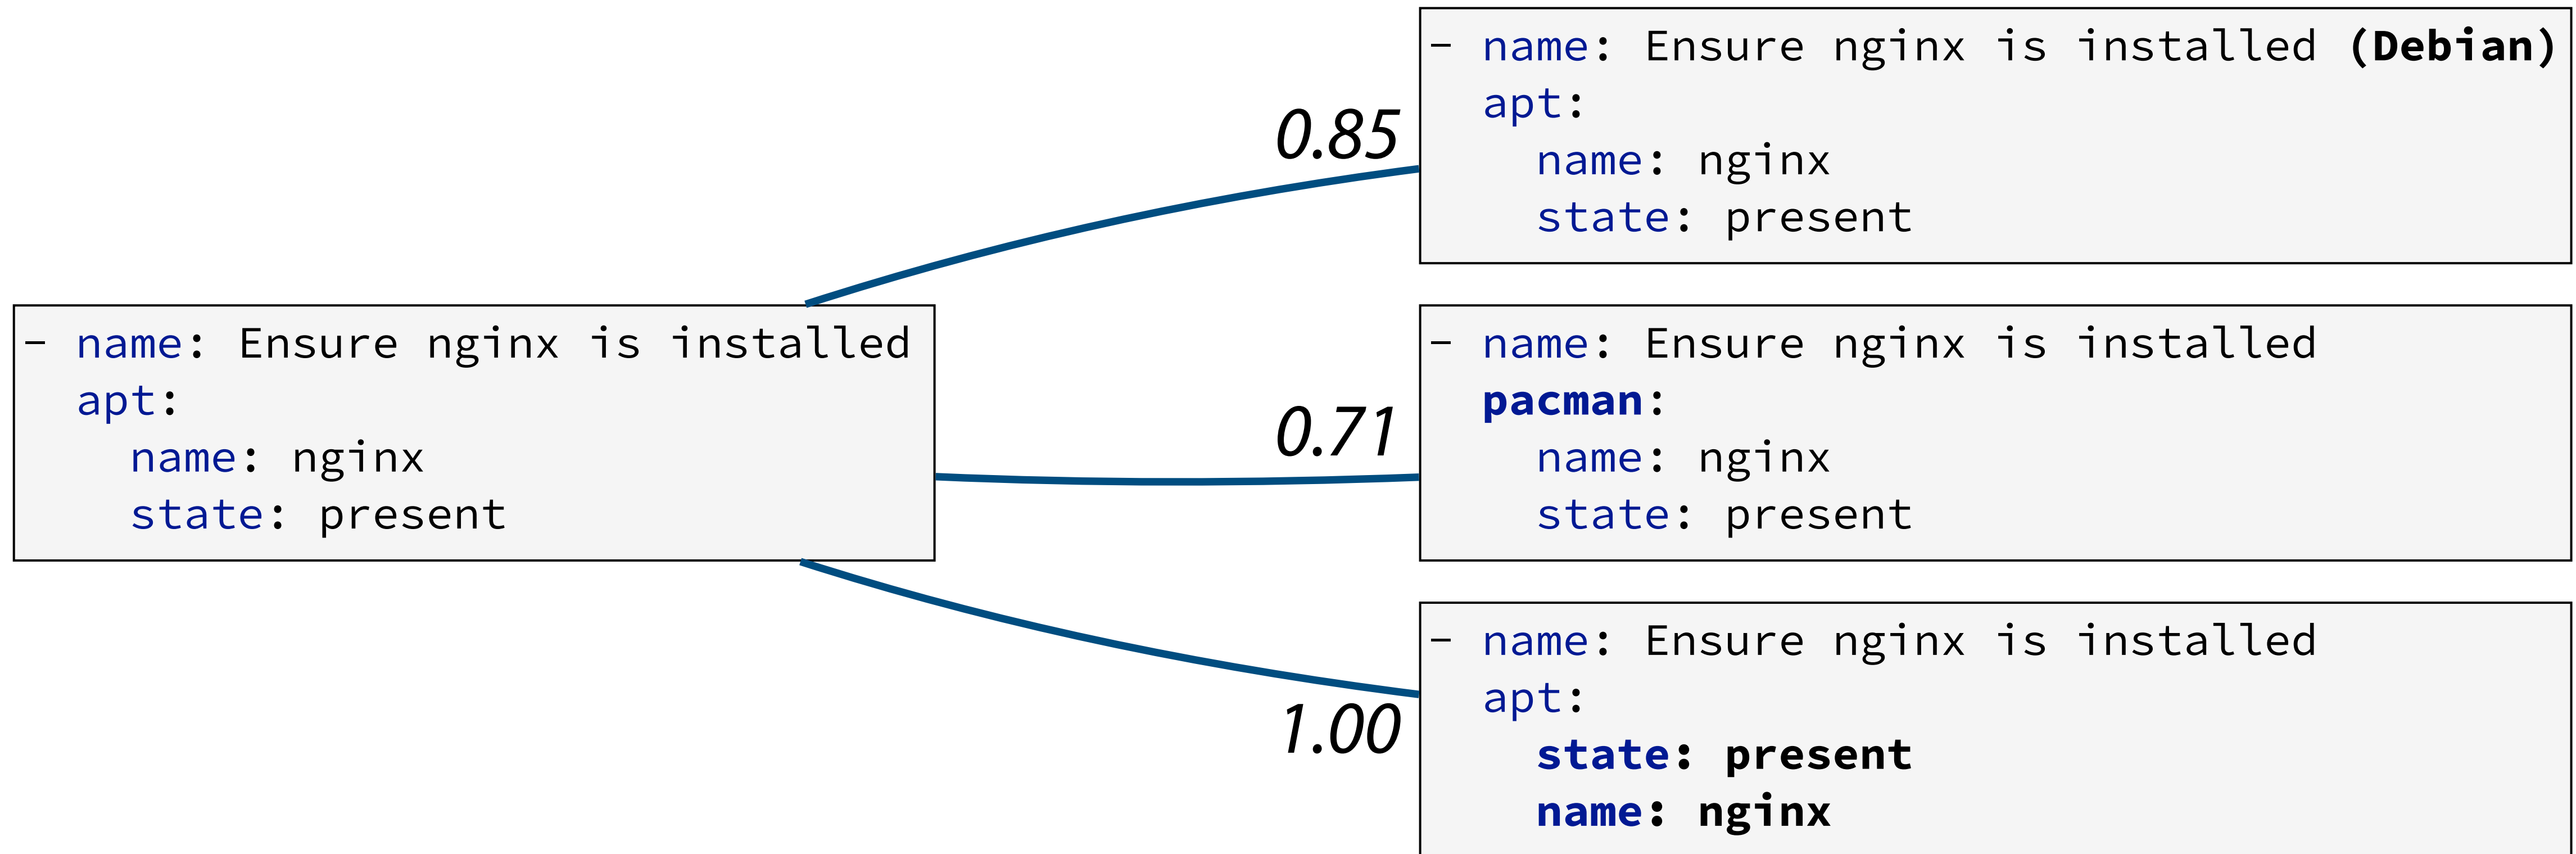


Local relocation



Too fine-grained

Task Similarity

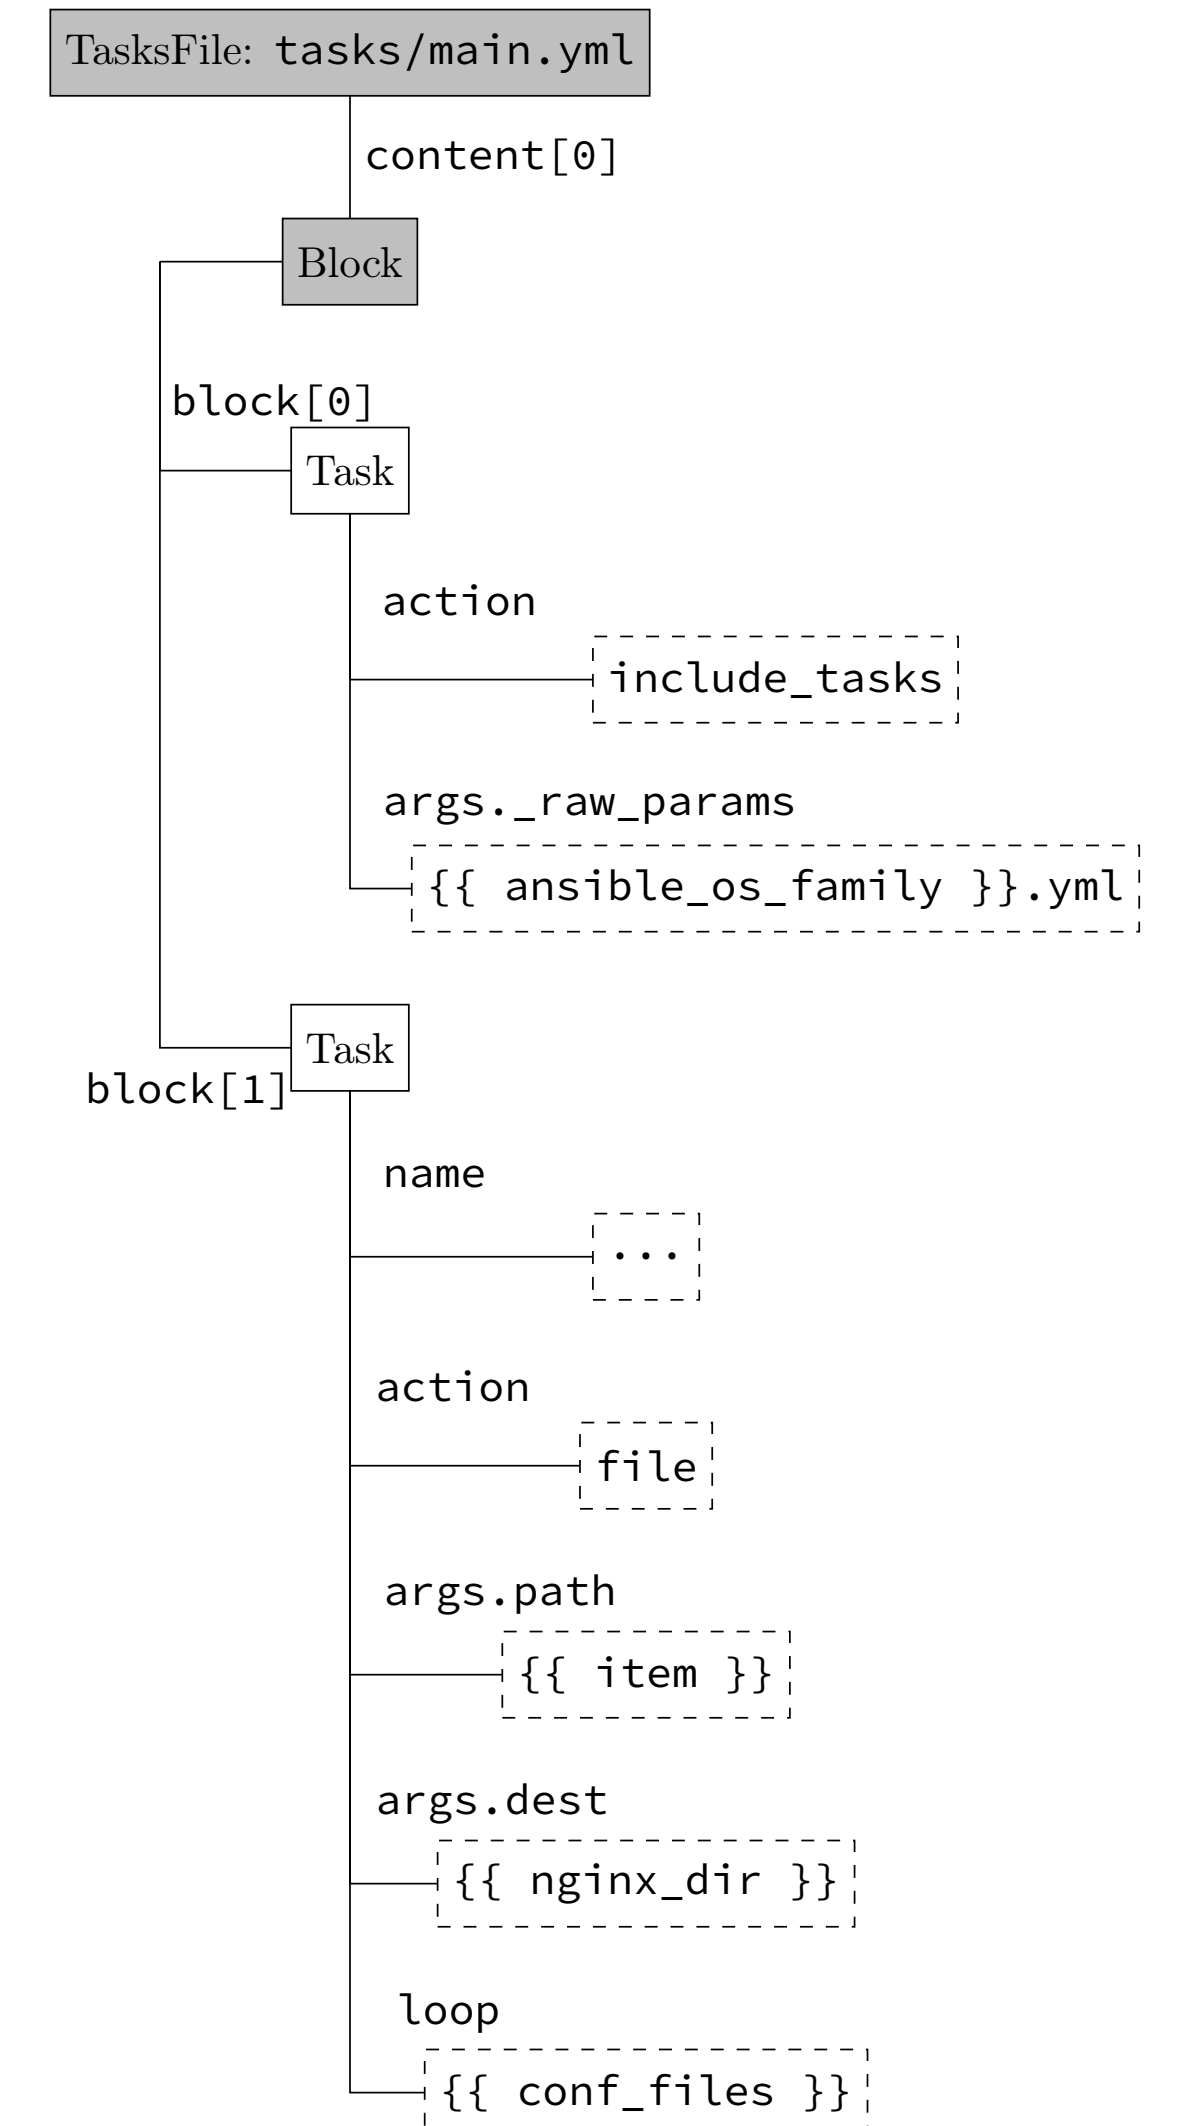
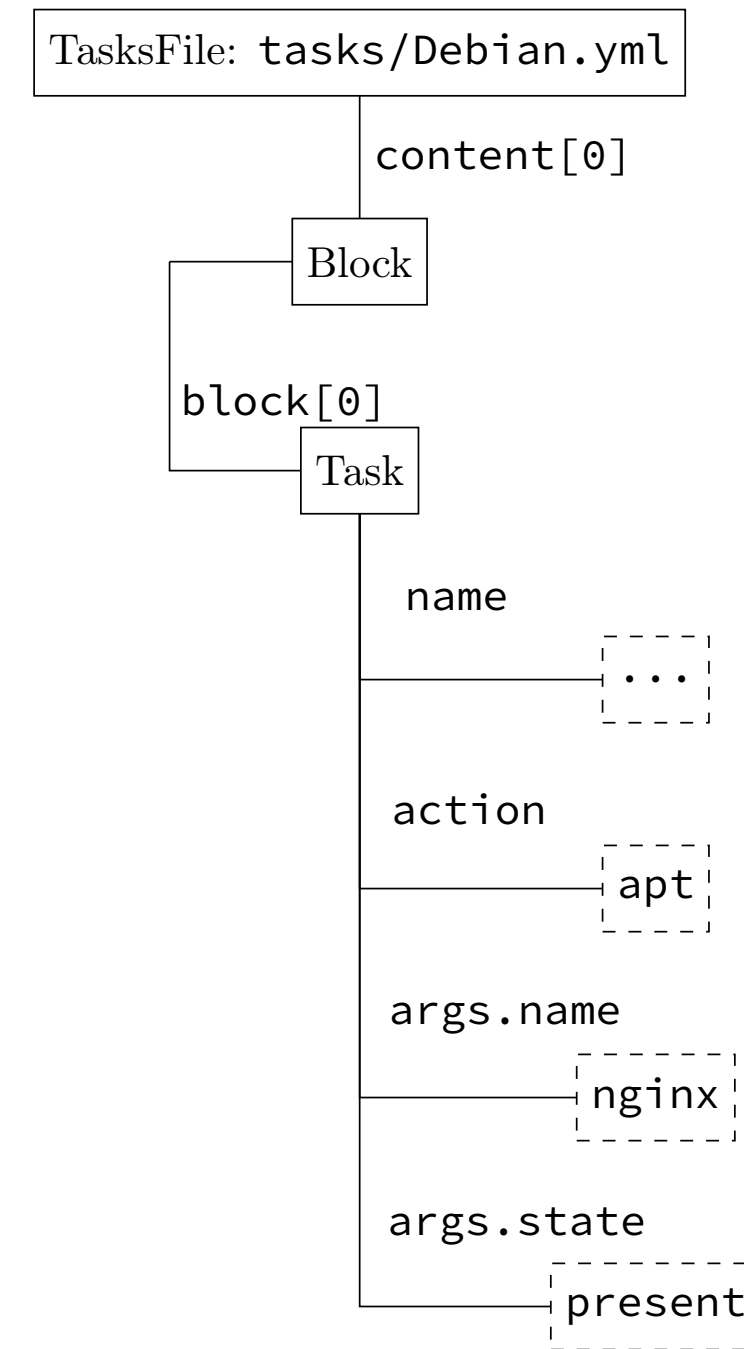
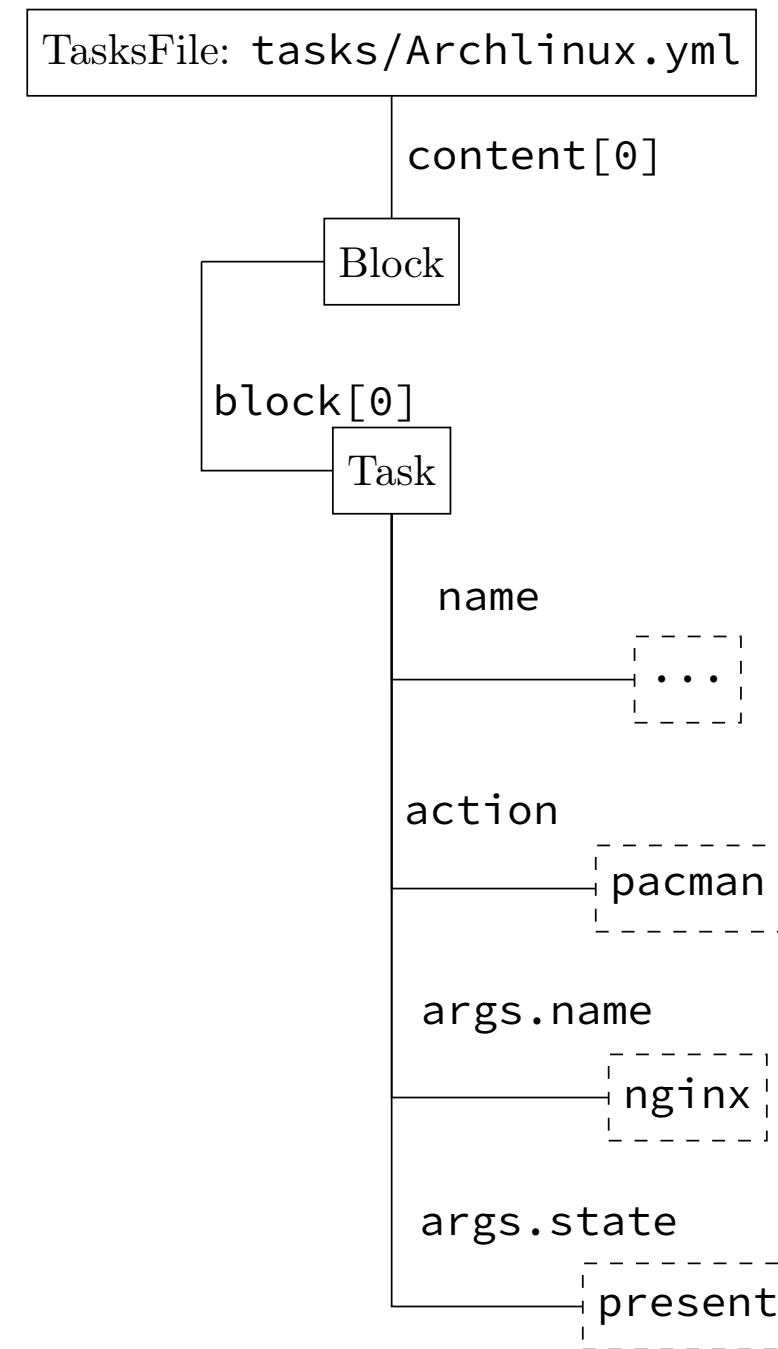
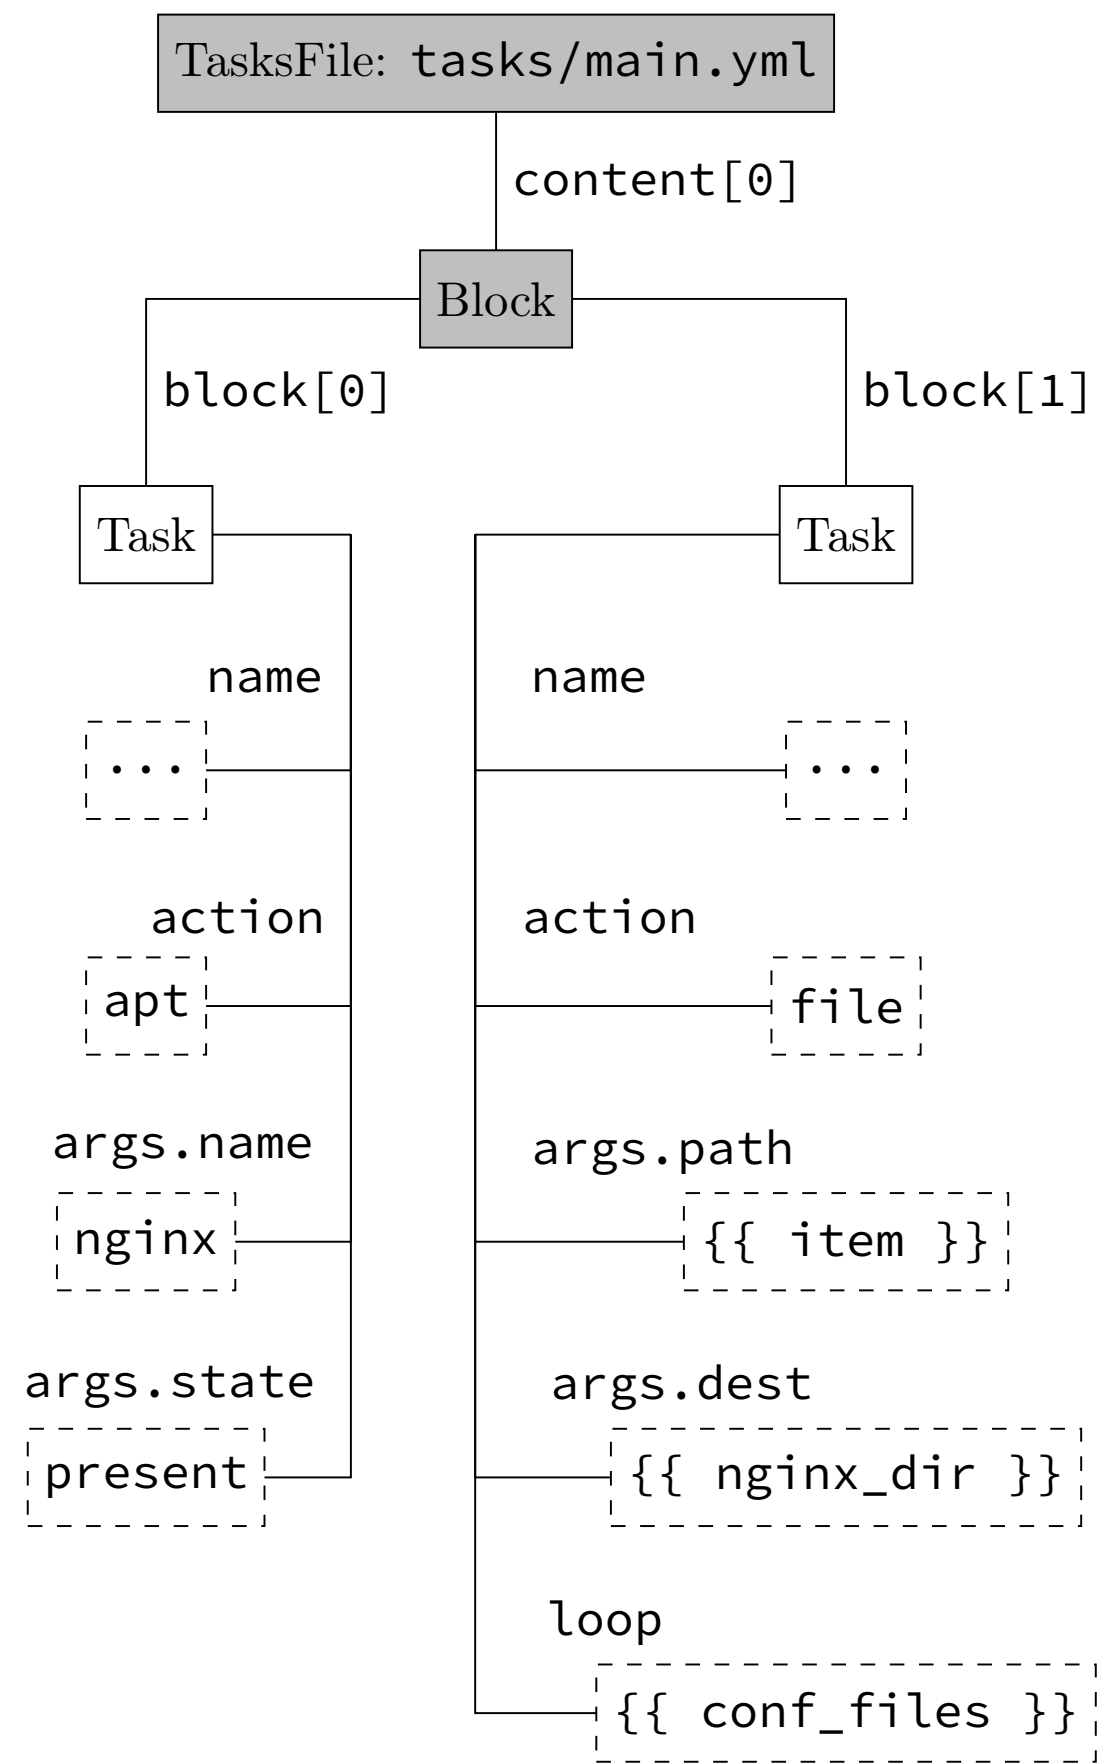


$$sim_{task}(T_1, T_2) = \frac{\sum_{kw \in T_1 \cap T_2} w(kw)}{\sum_{kw \in T_1 \cup T_2} w(kw)}$$

$$w(kw) = \begin{cases} 1, & \text{if } kw \in \{\text{action, args, when, loop, loop_control}\} \\ 0.5, & \text{otherwise} \end{cases}$$

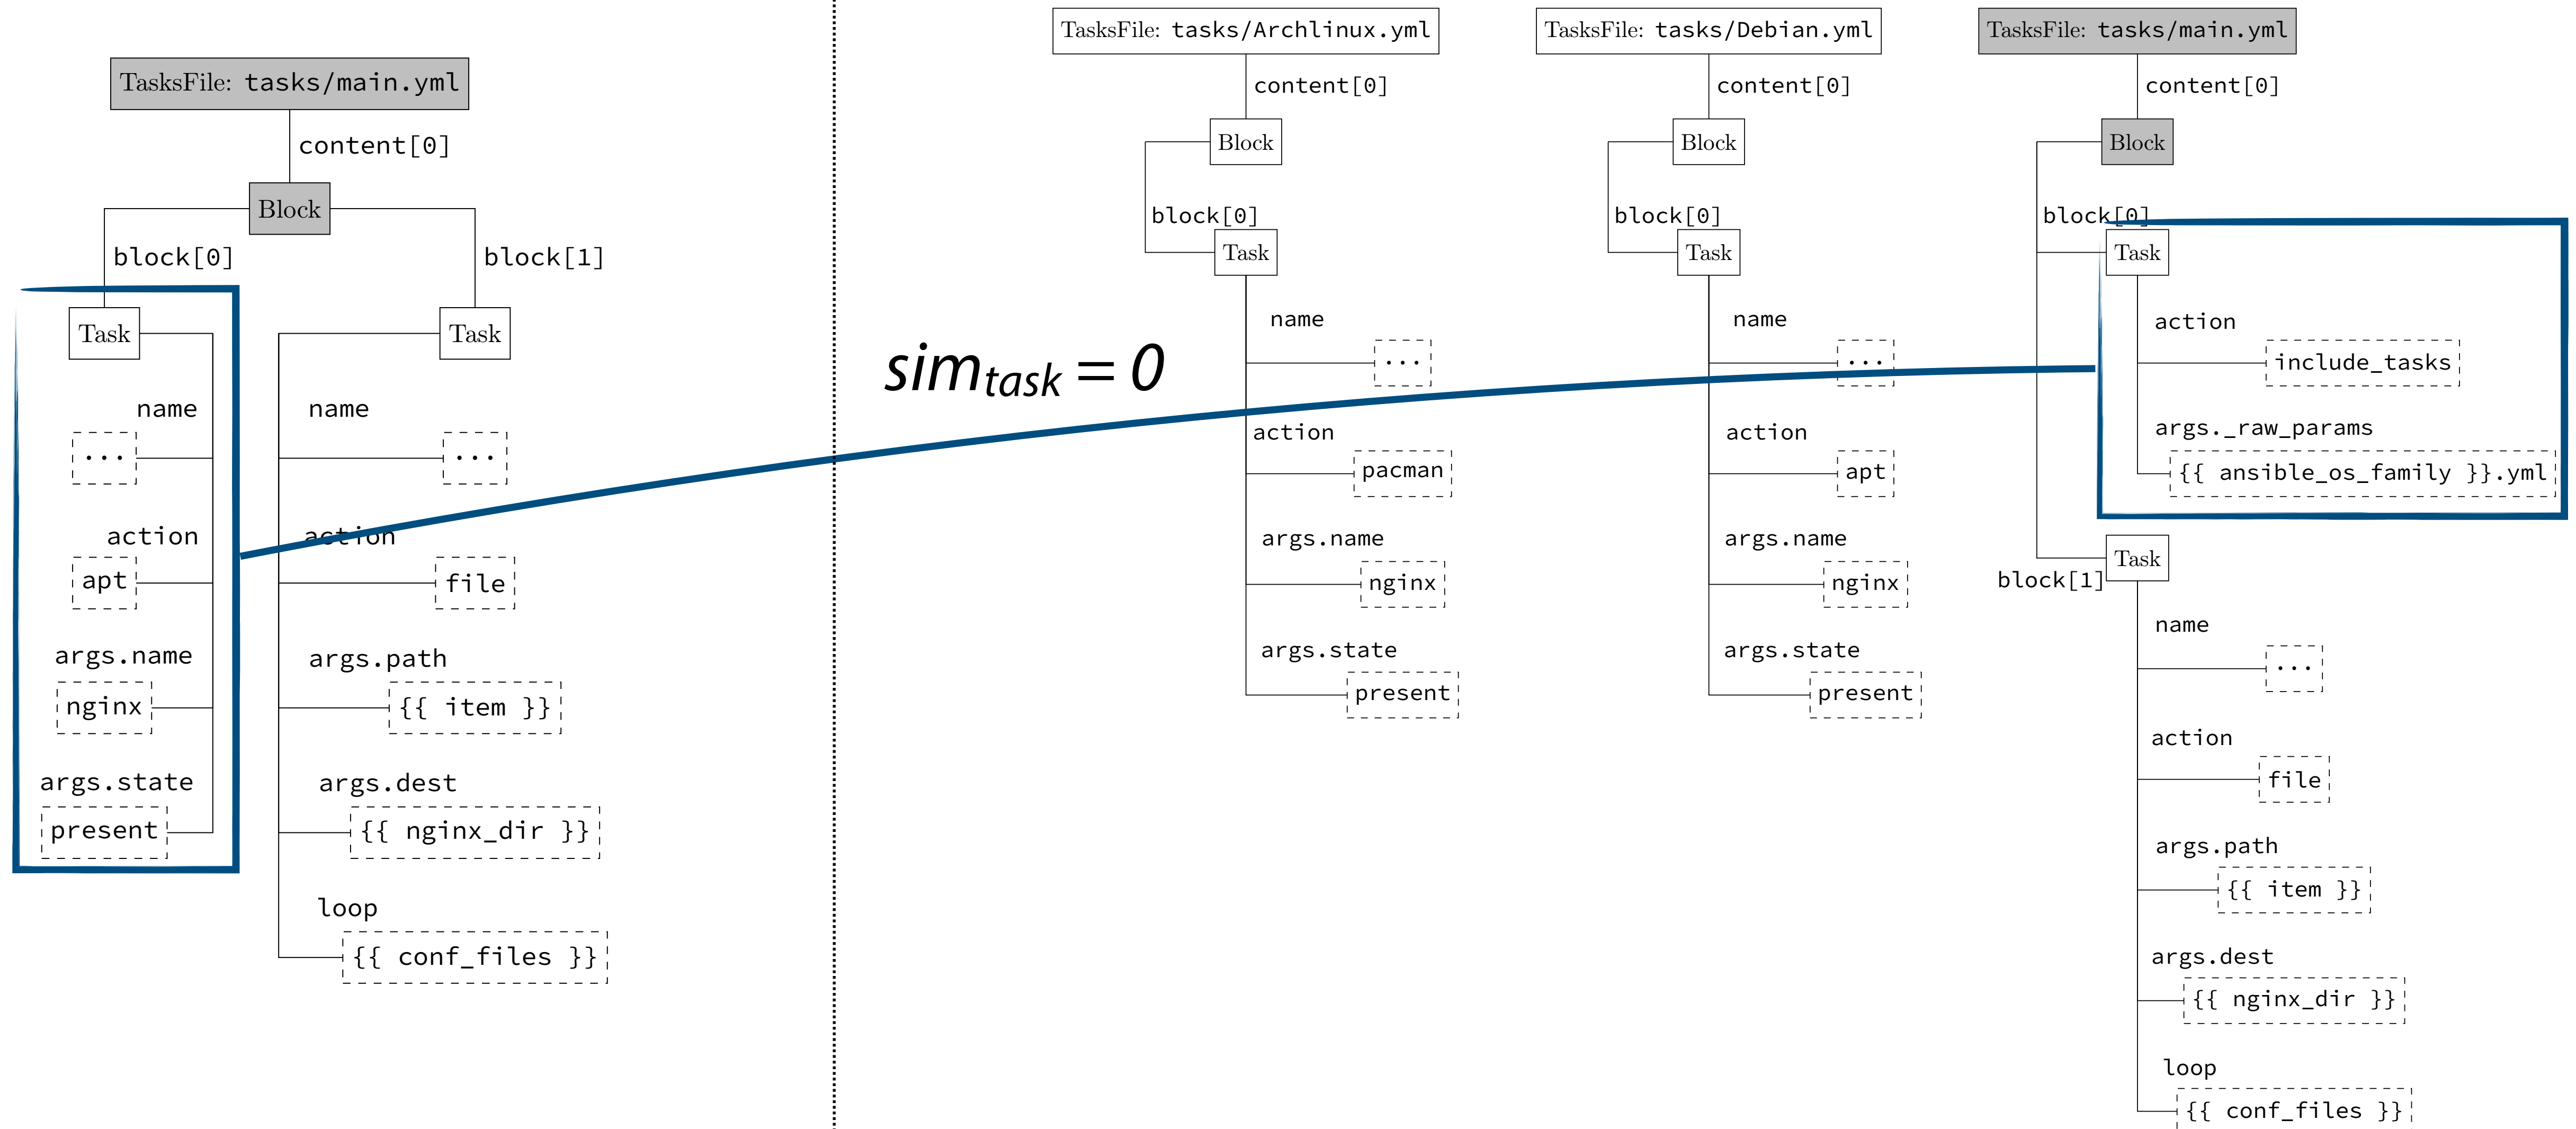
“Weighted Jaccard”

Task Matching



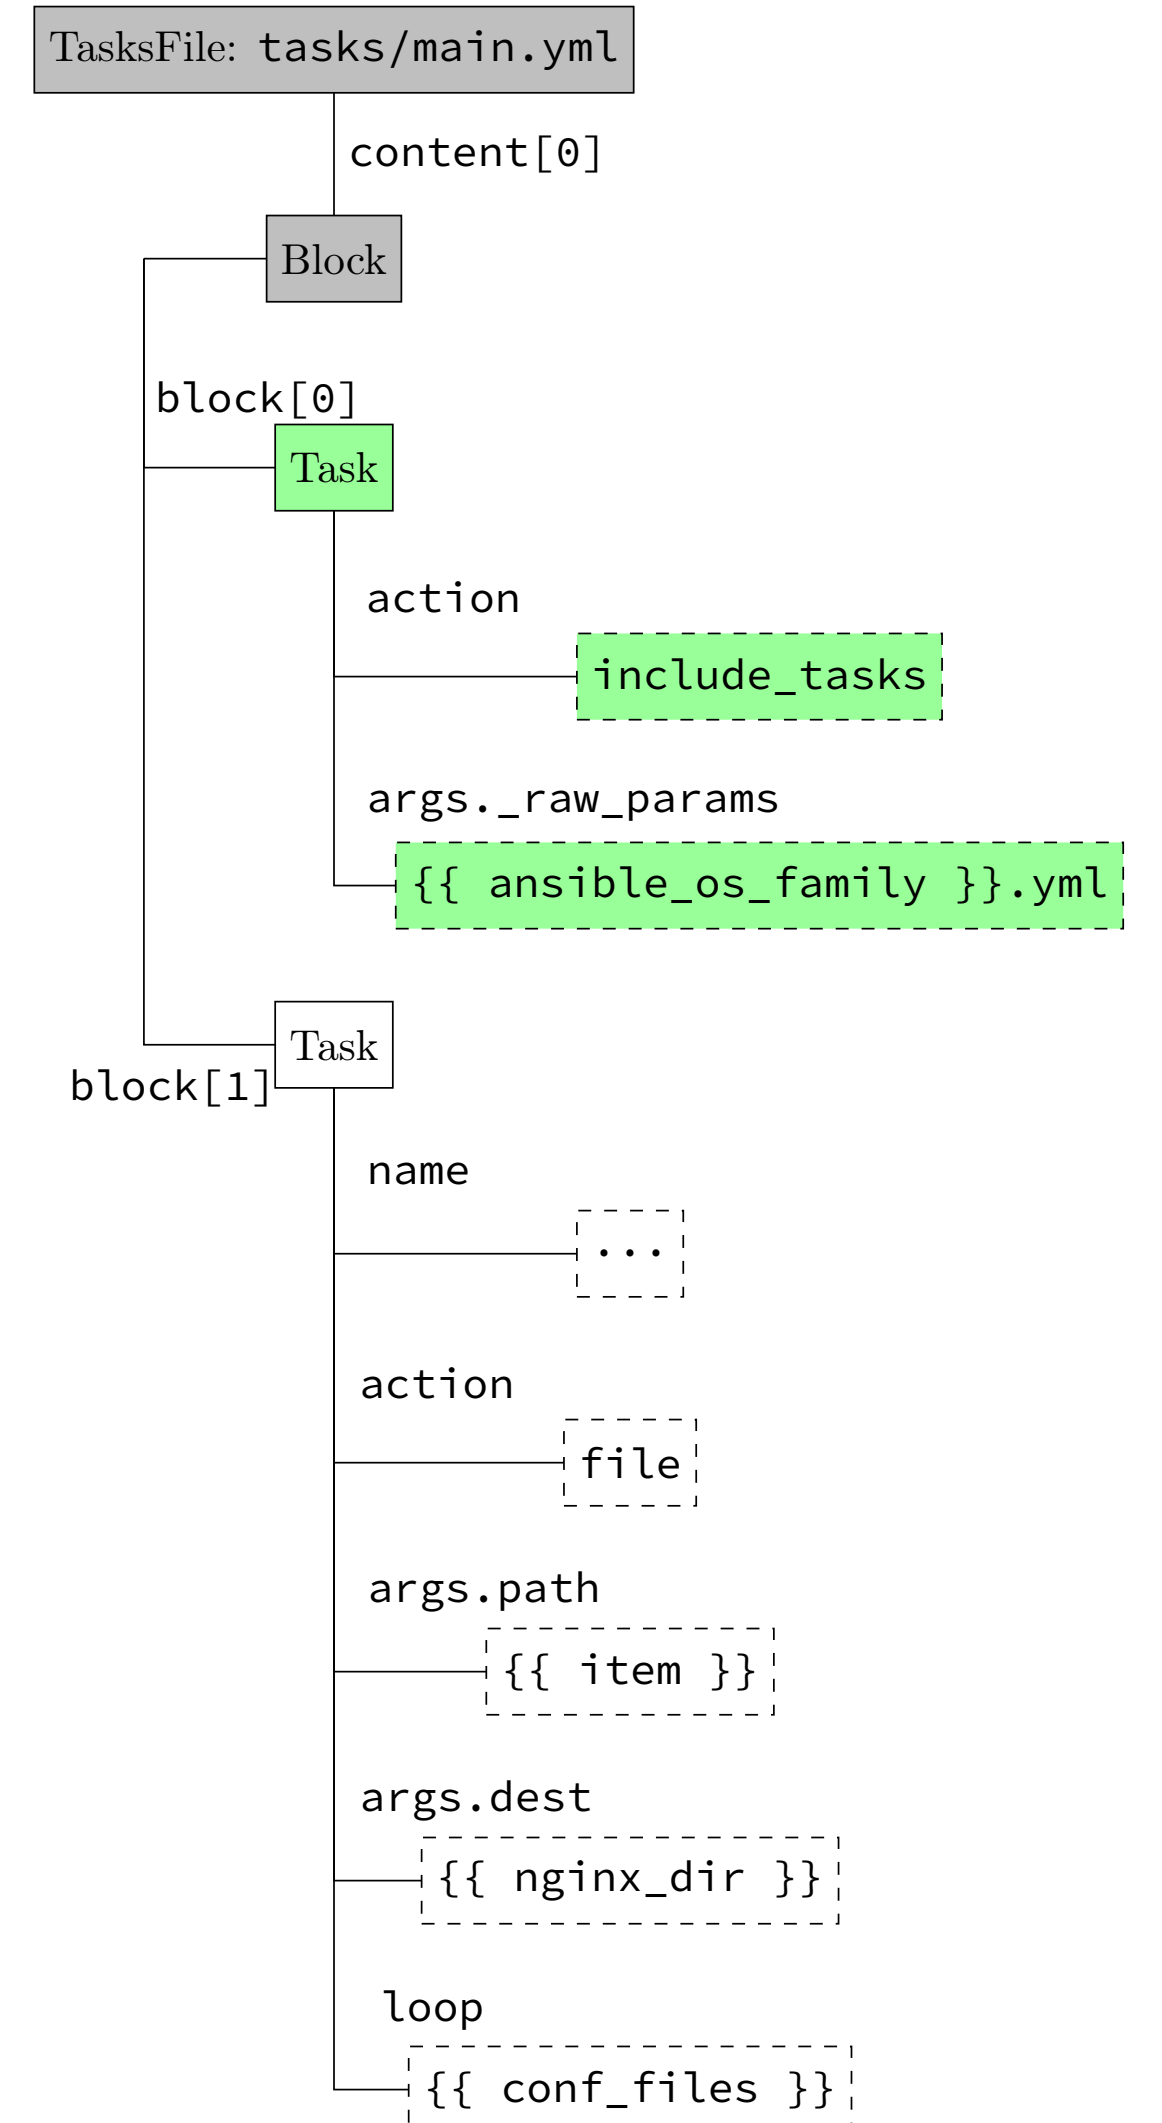
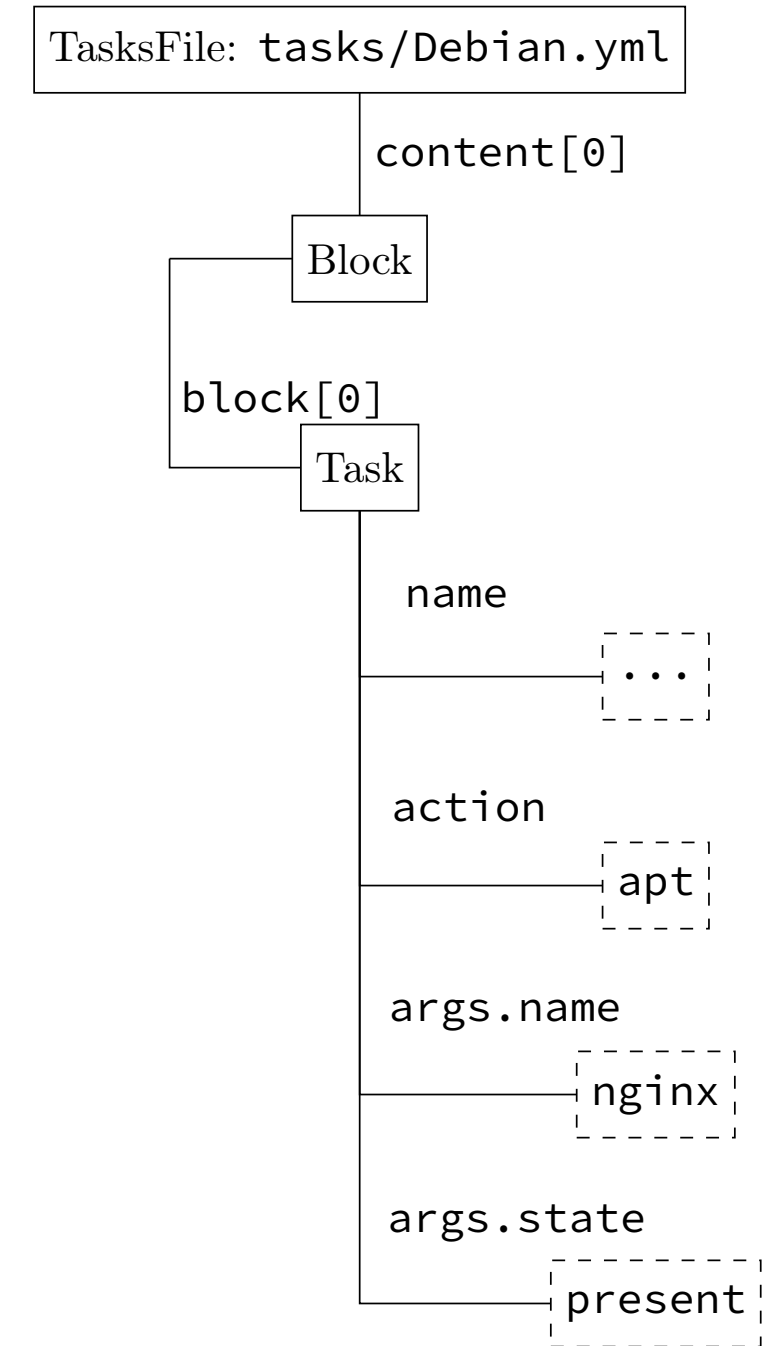
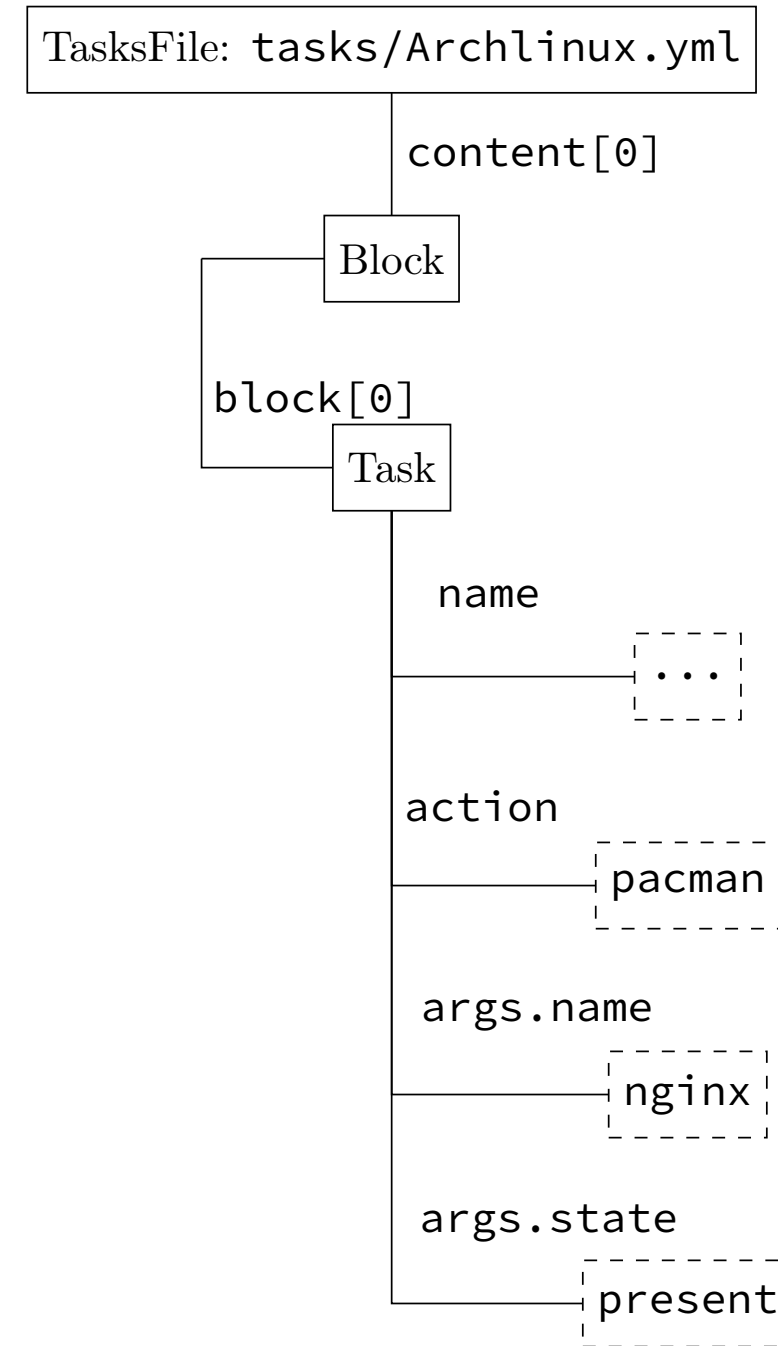
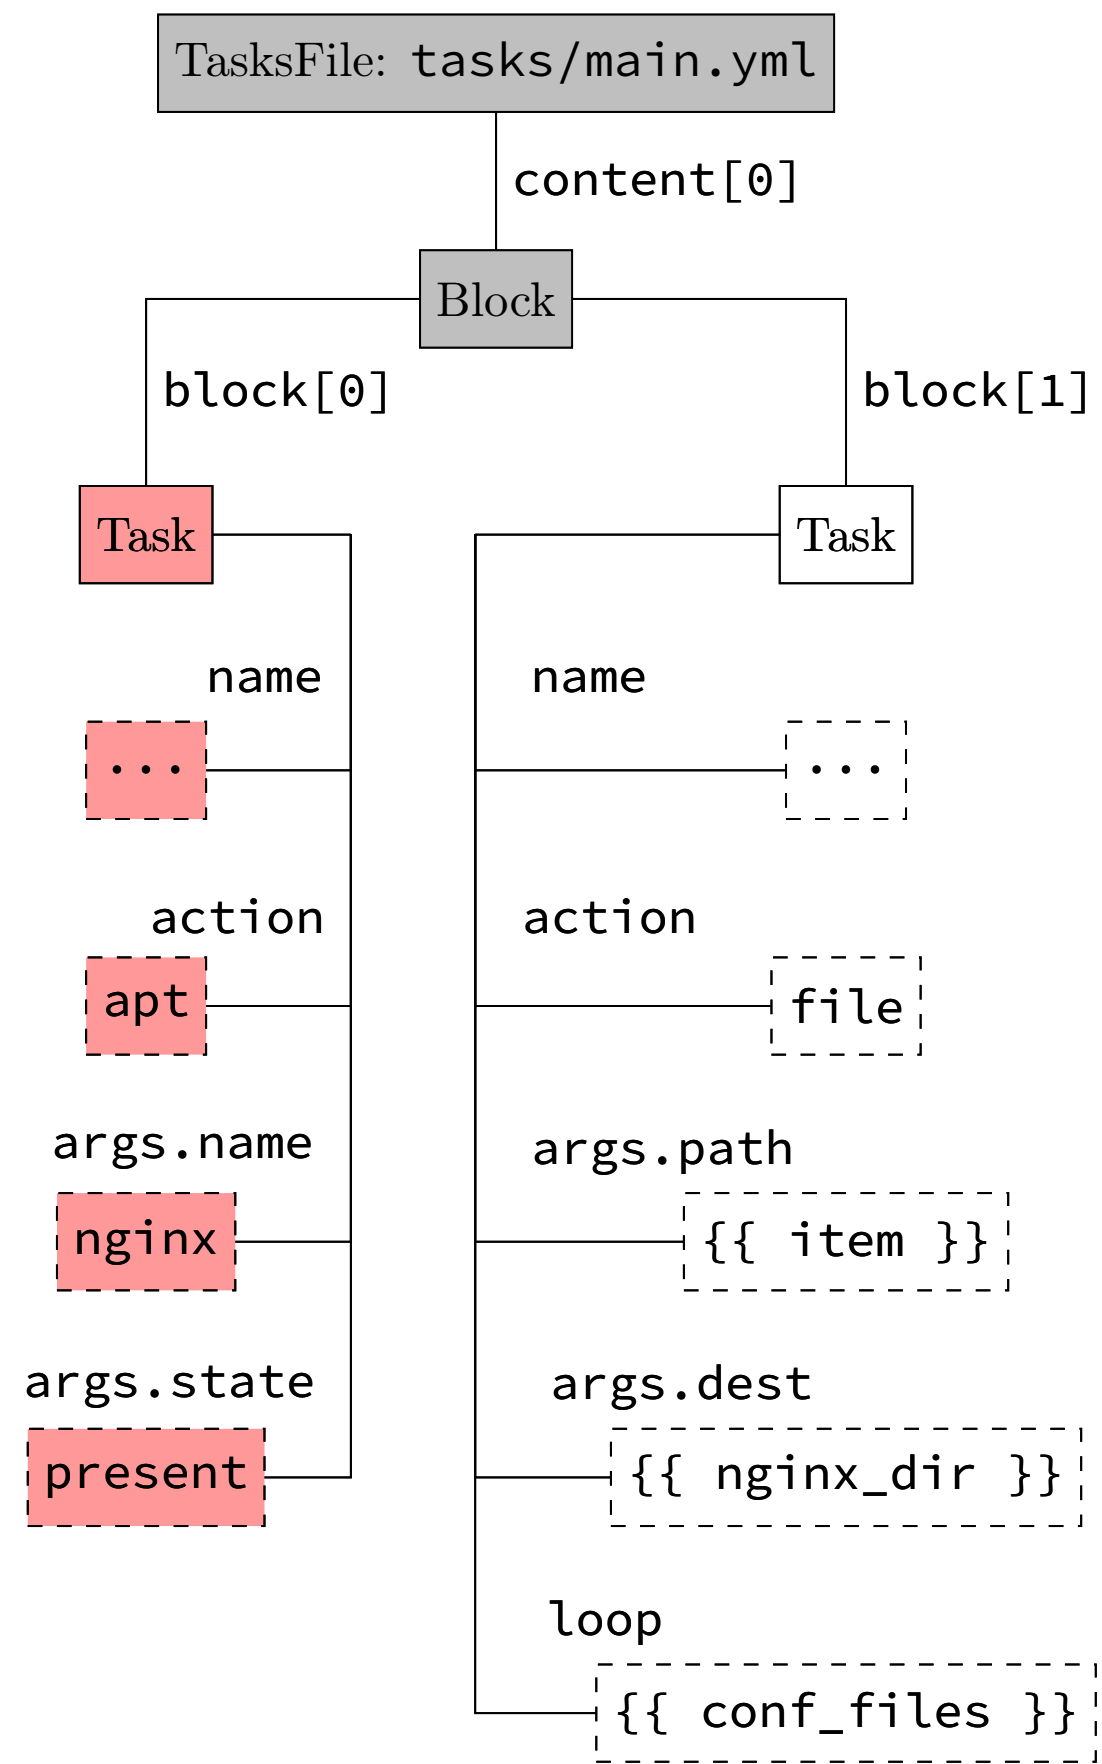
v1 v2

Task Matching



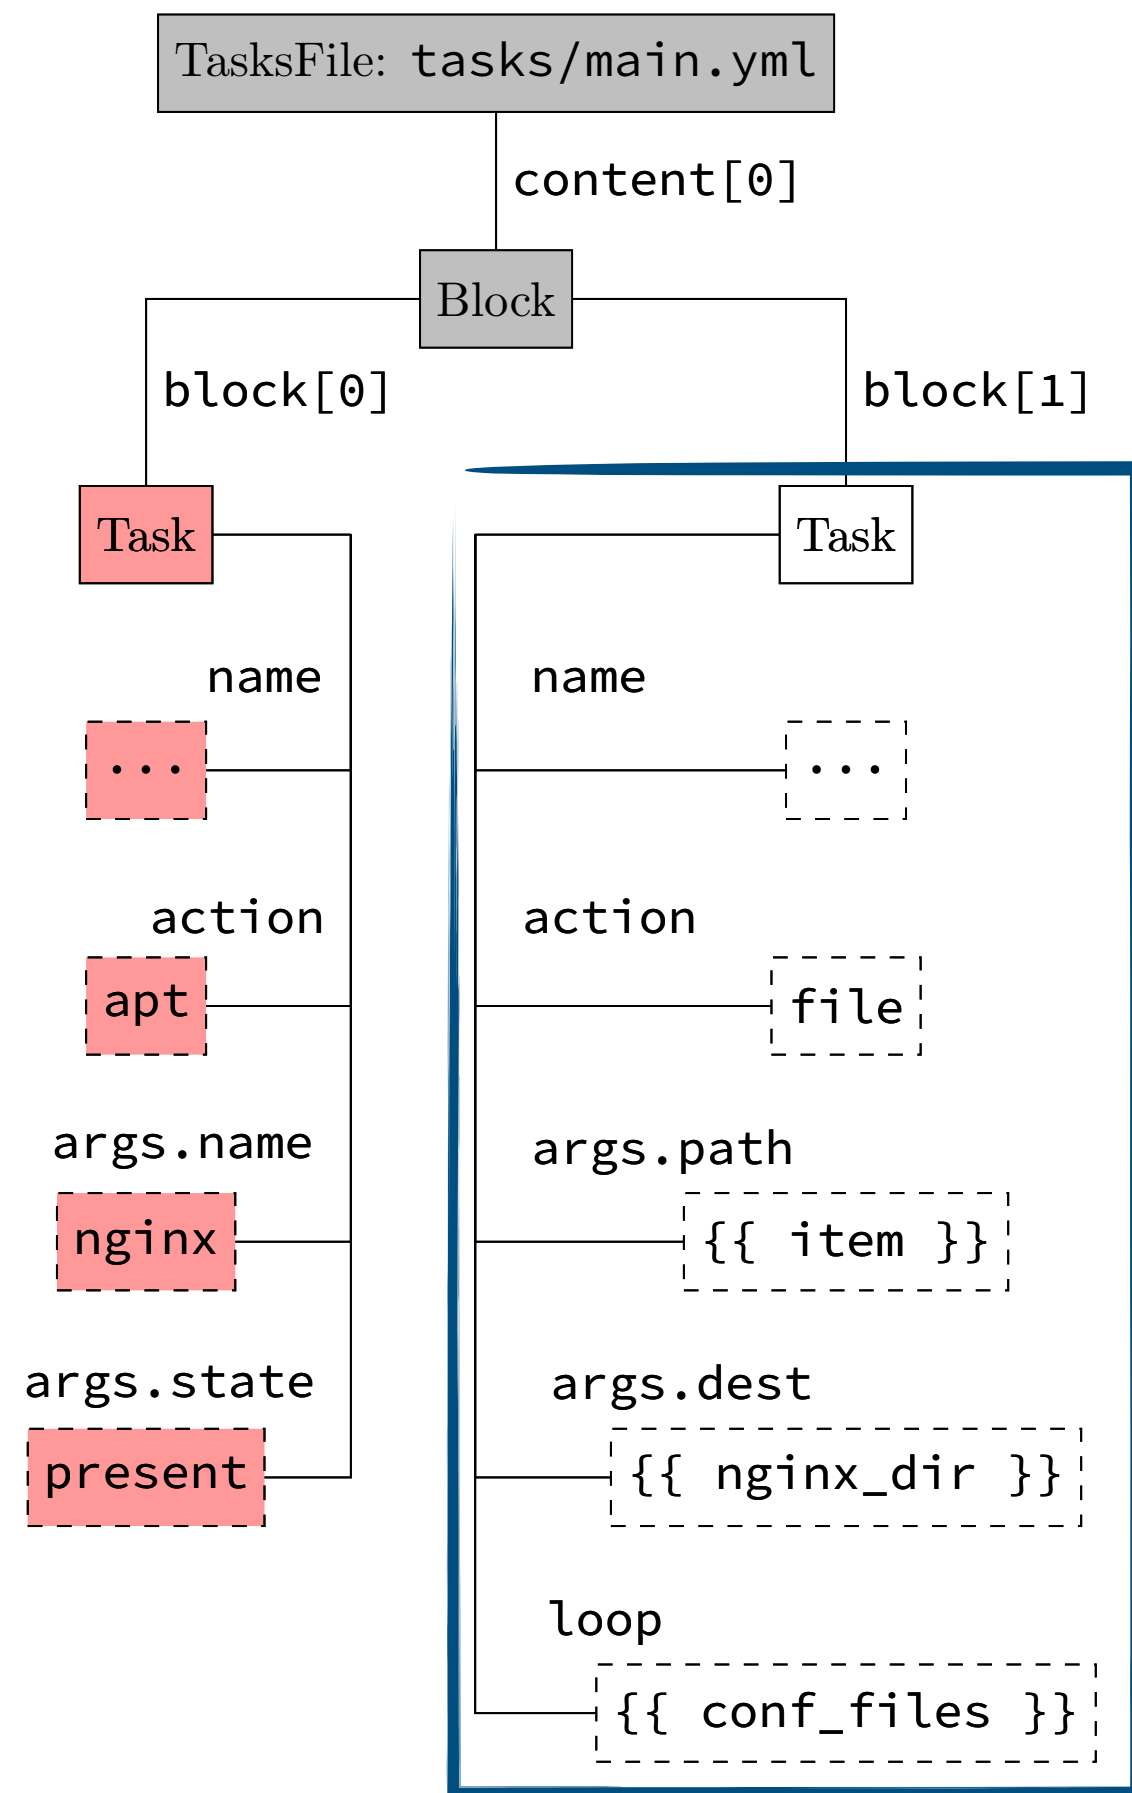
v1 v2

Task Matching



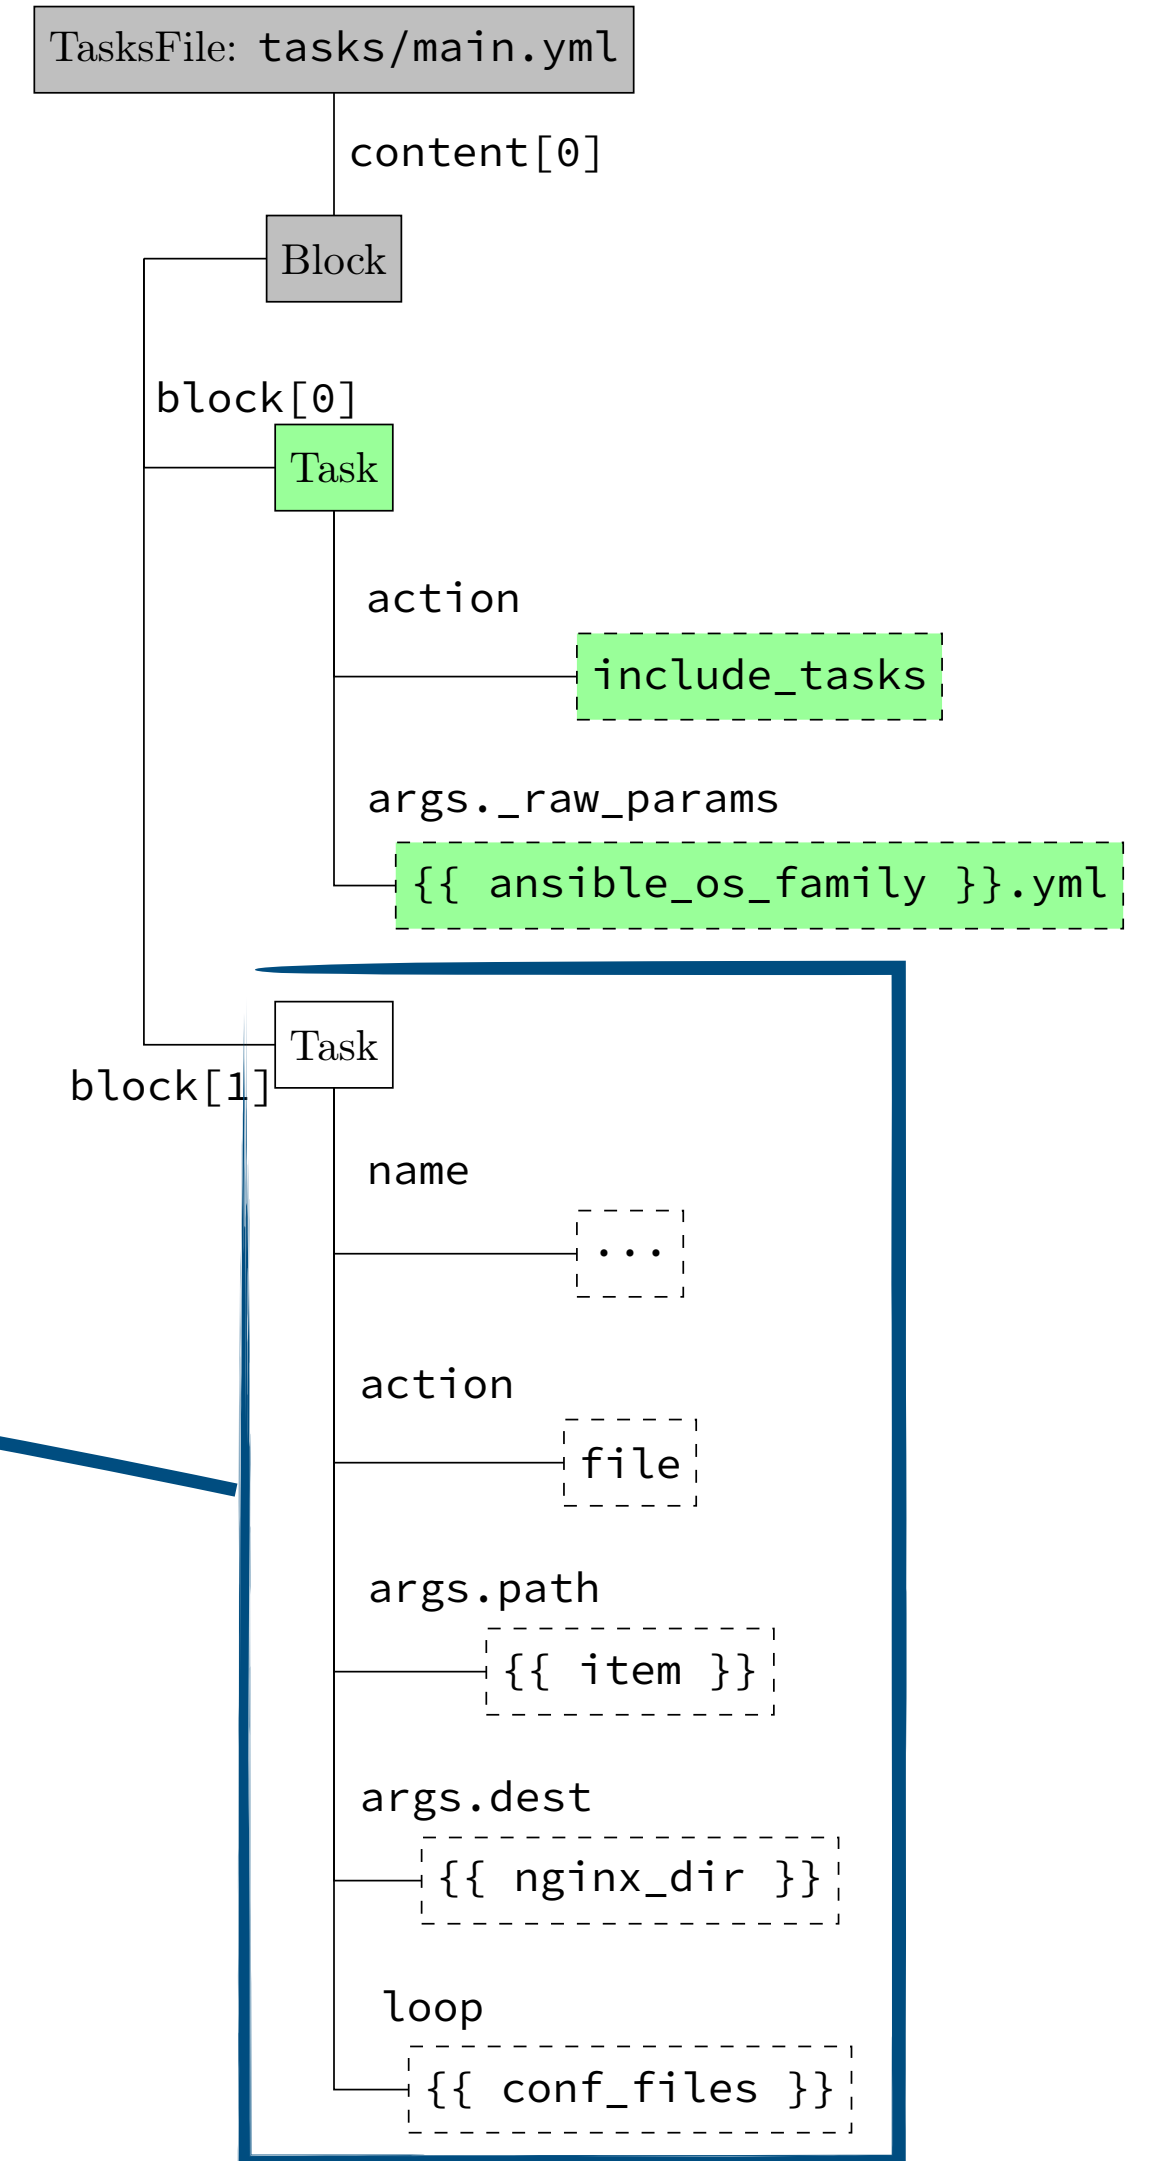
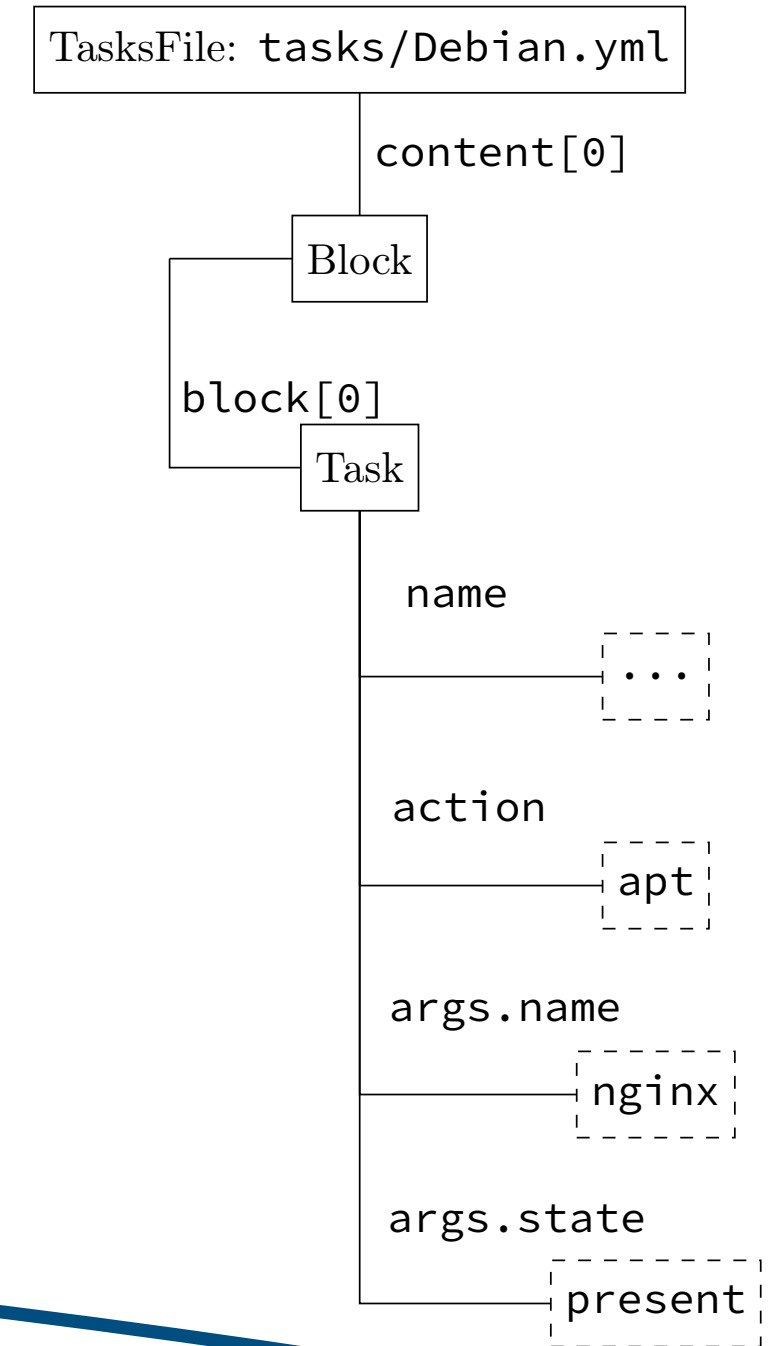
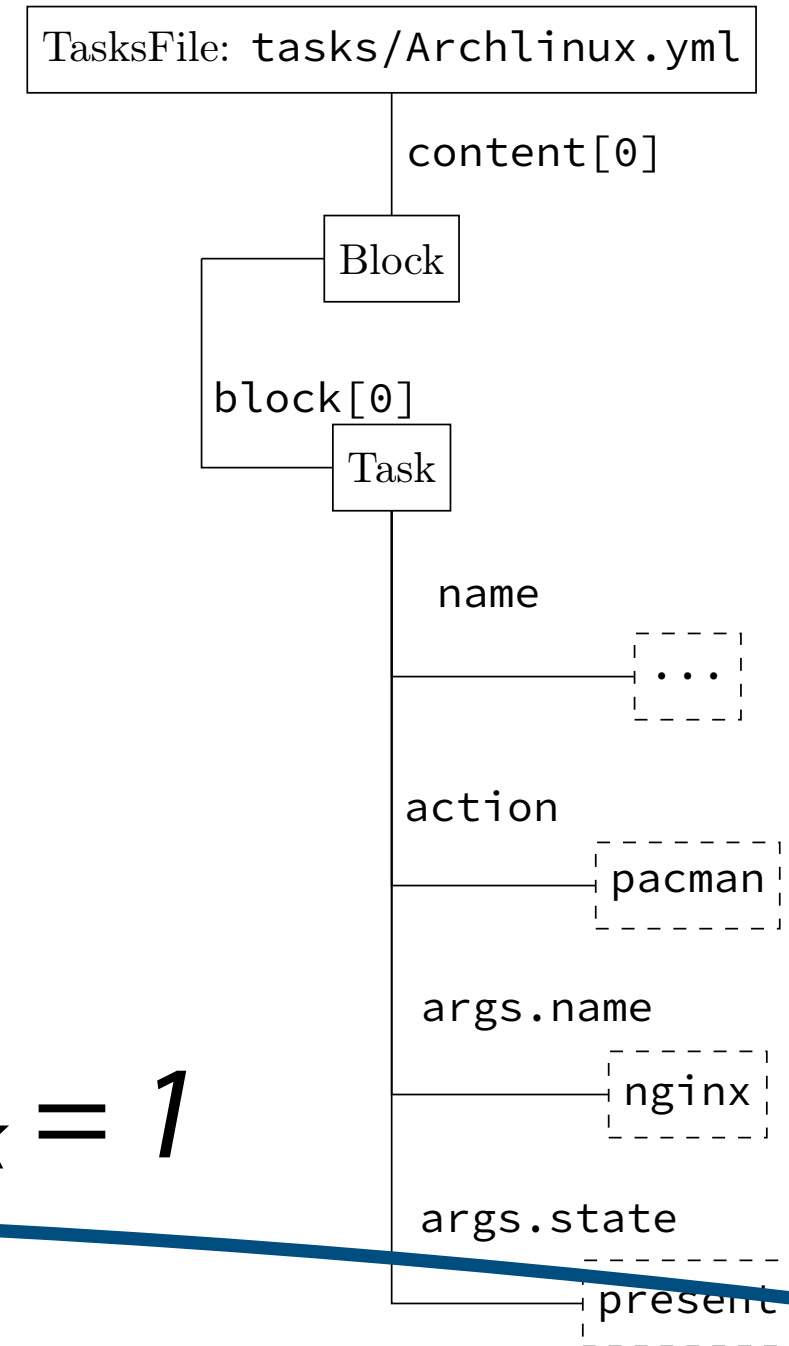
v1 v2

Task Matching

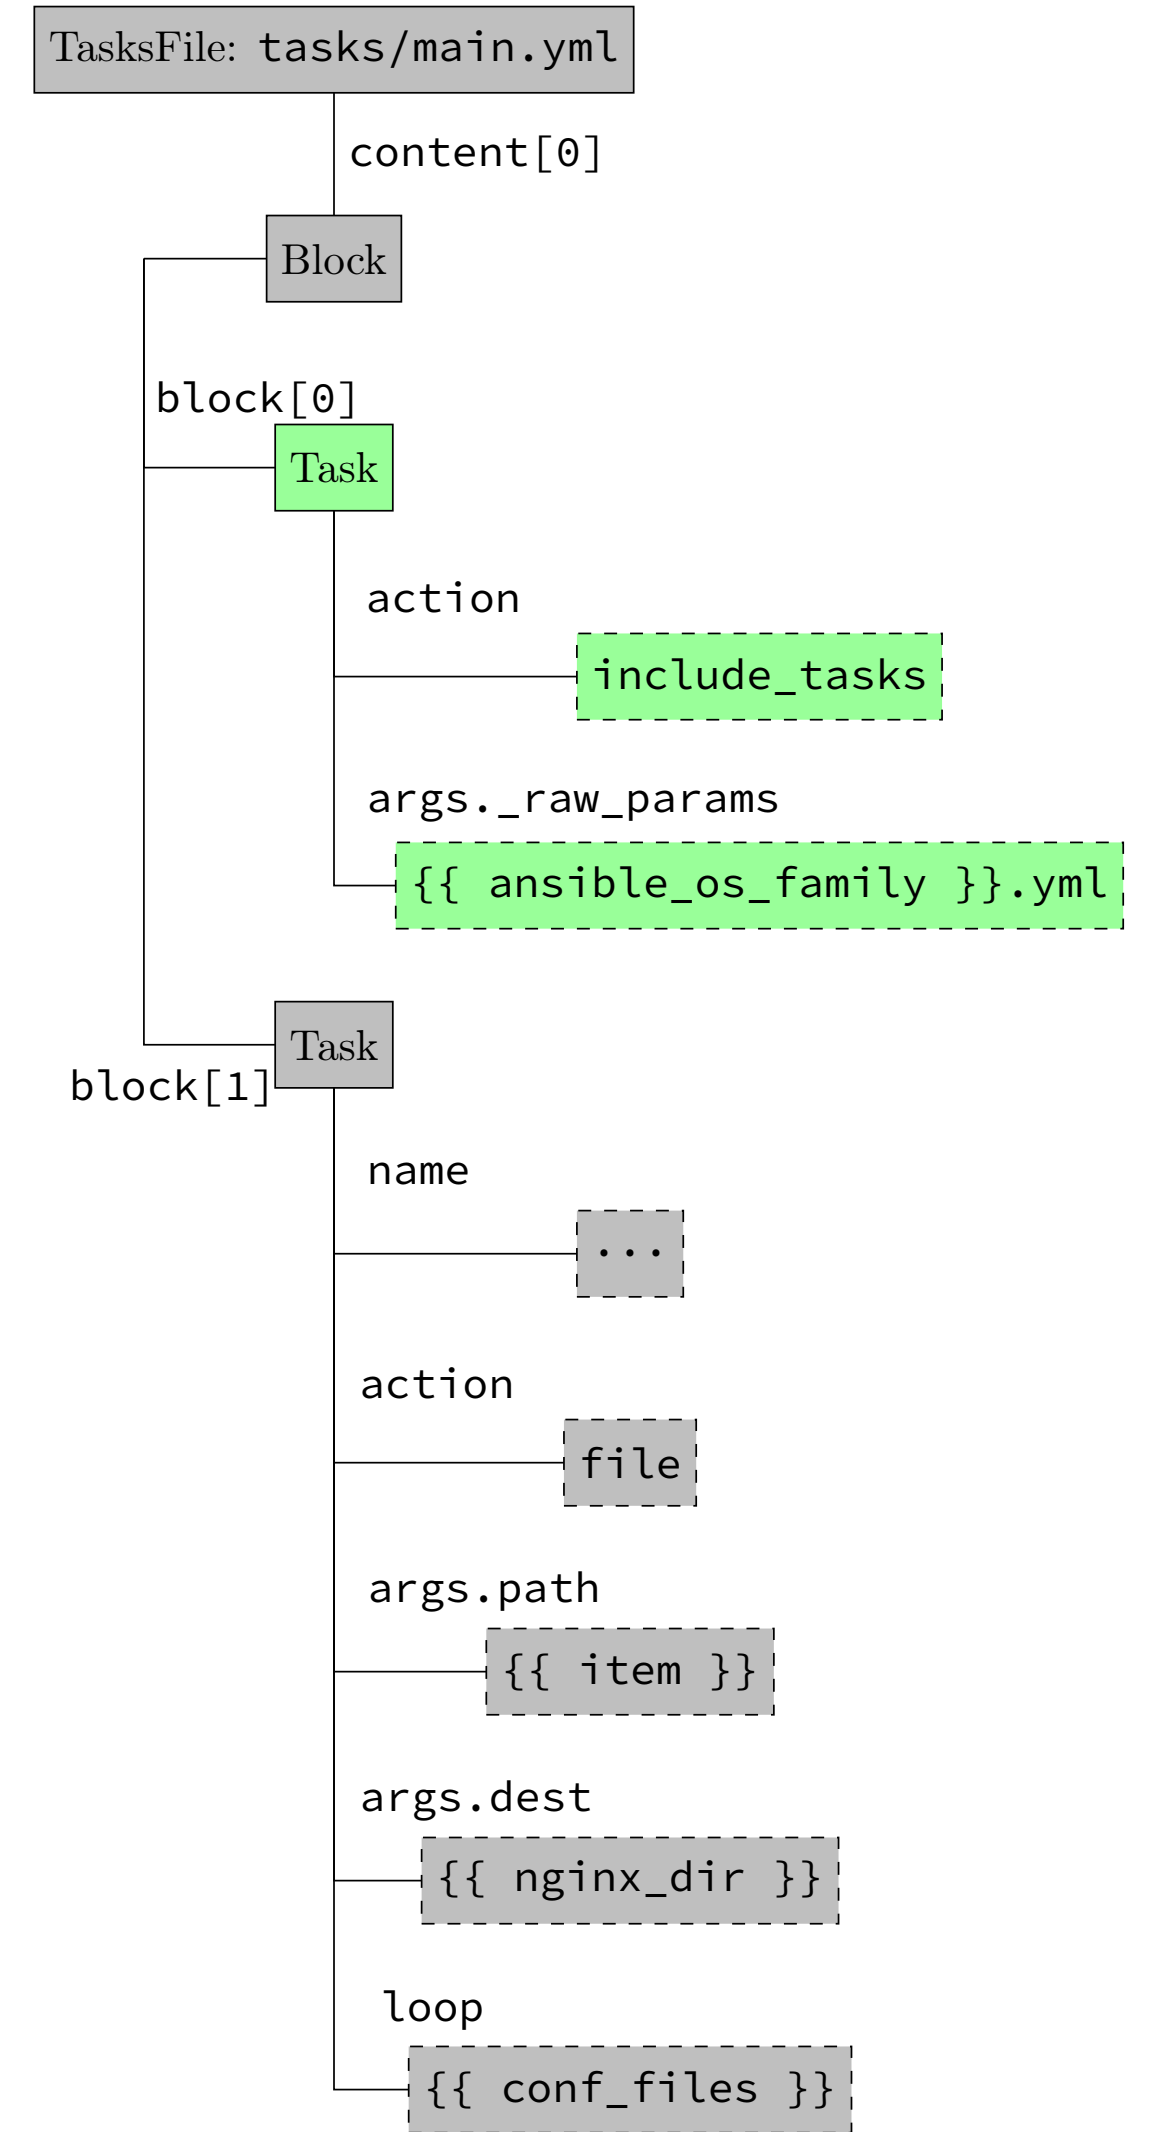
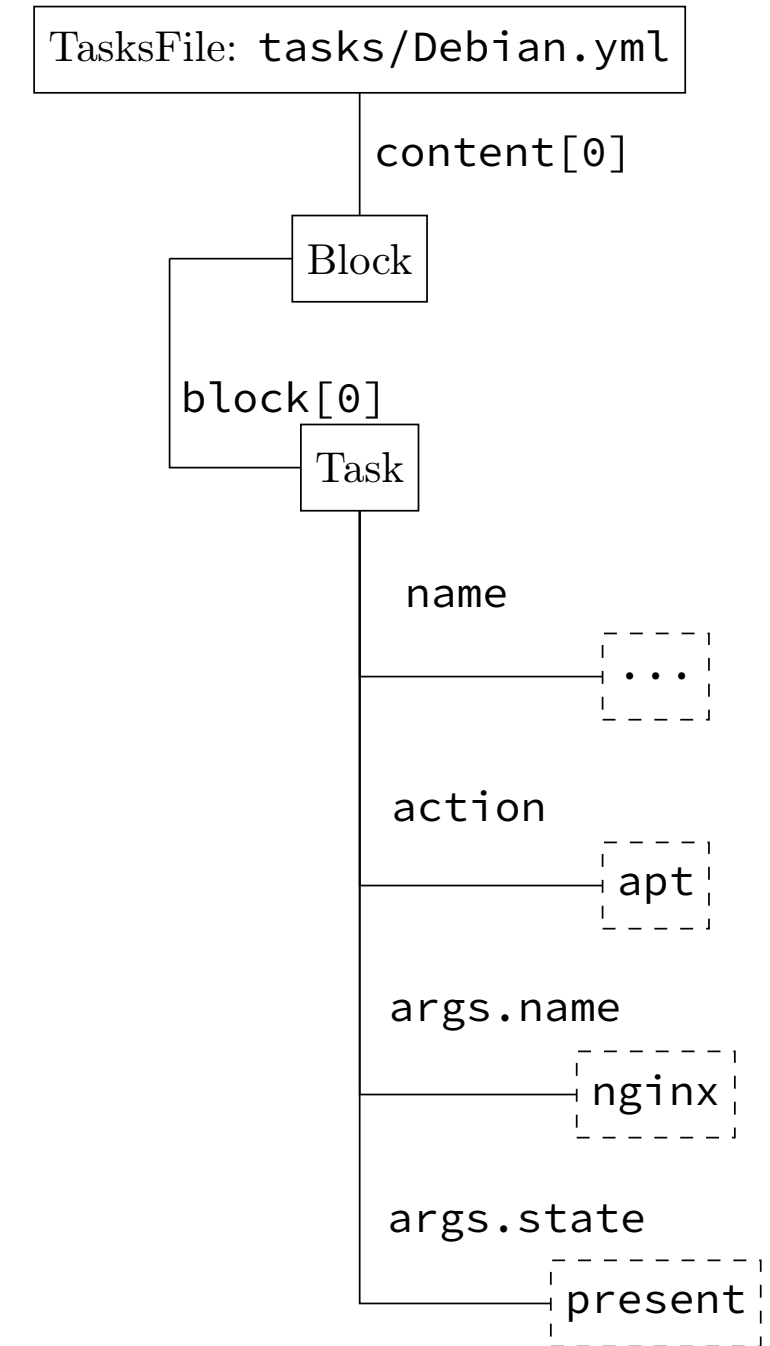
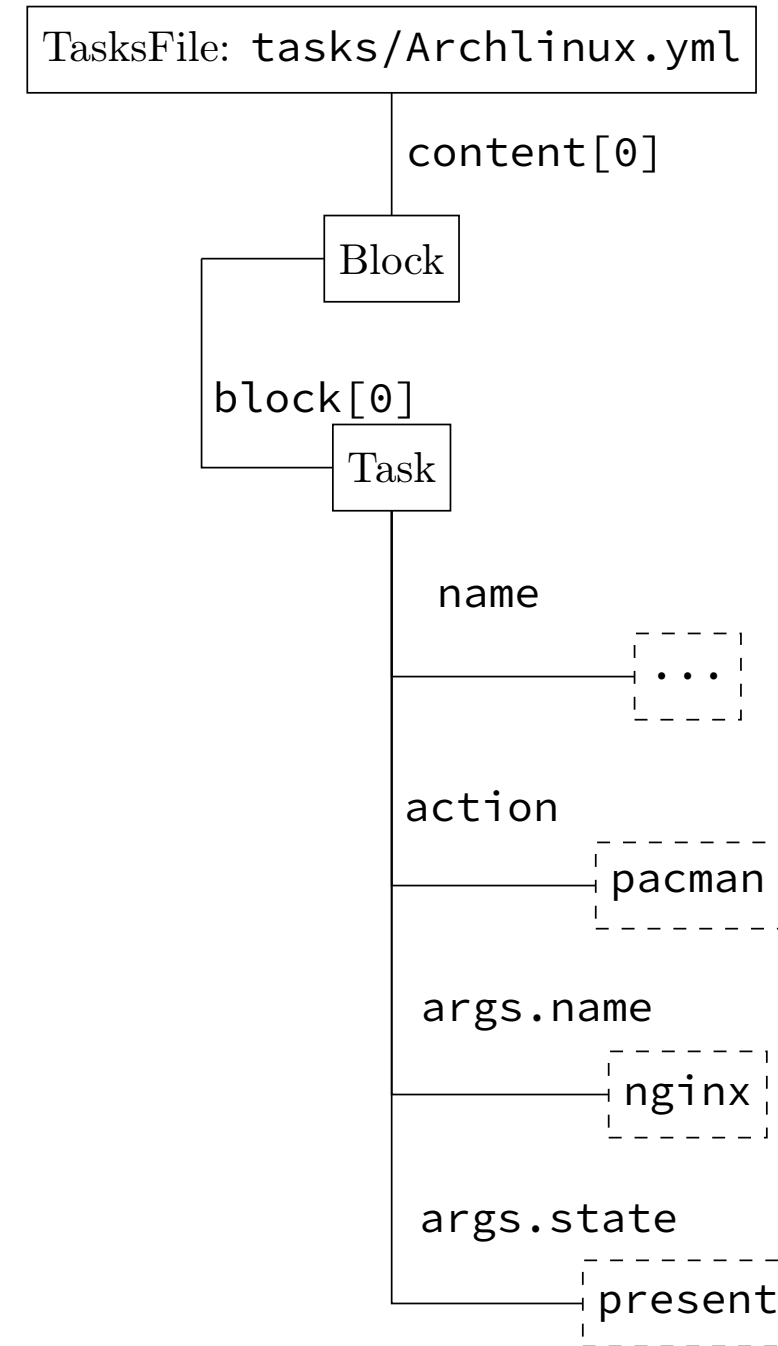
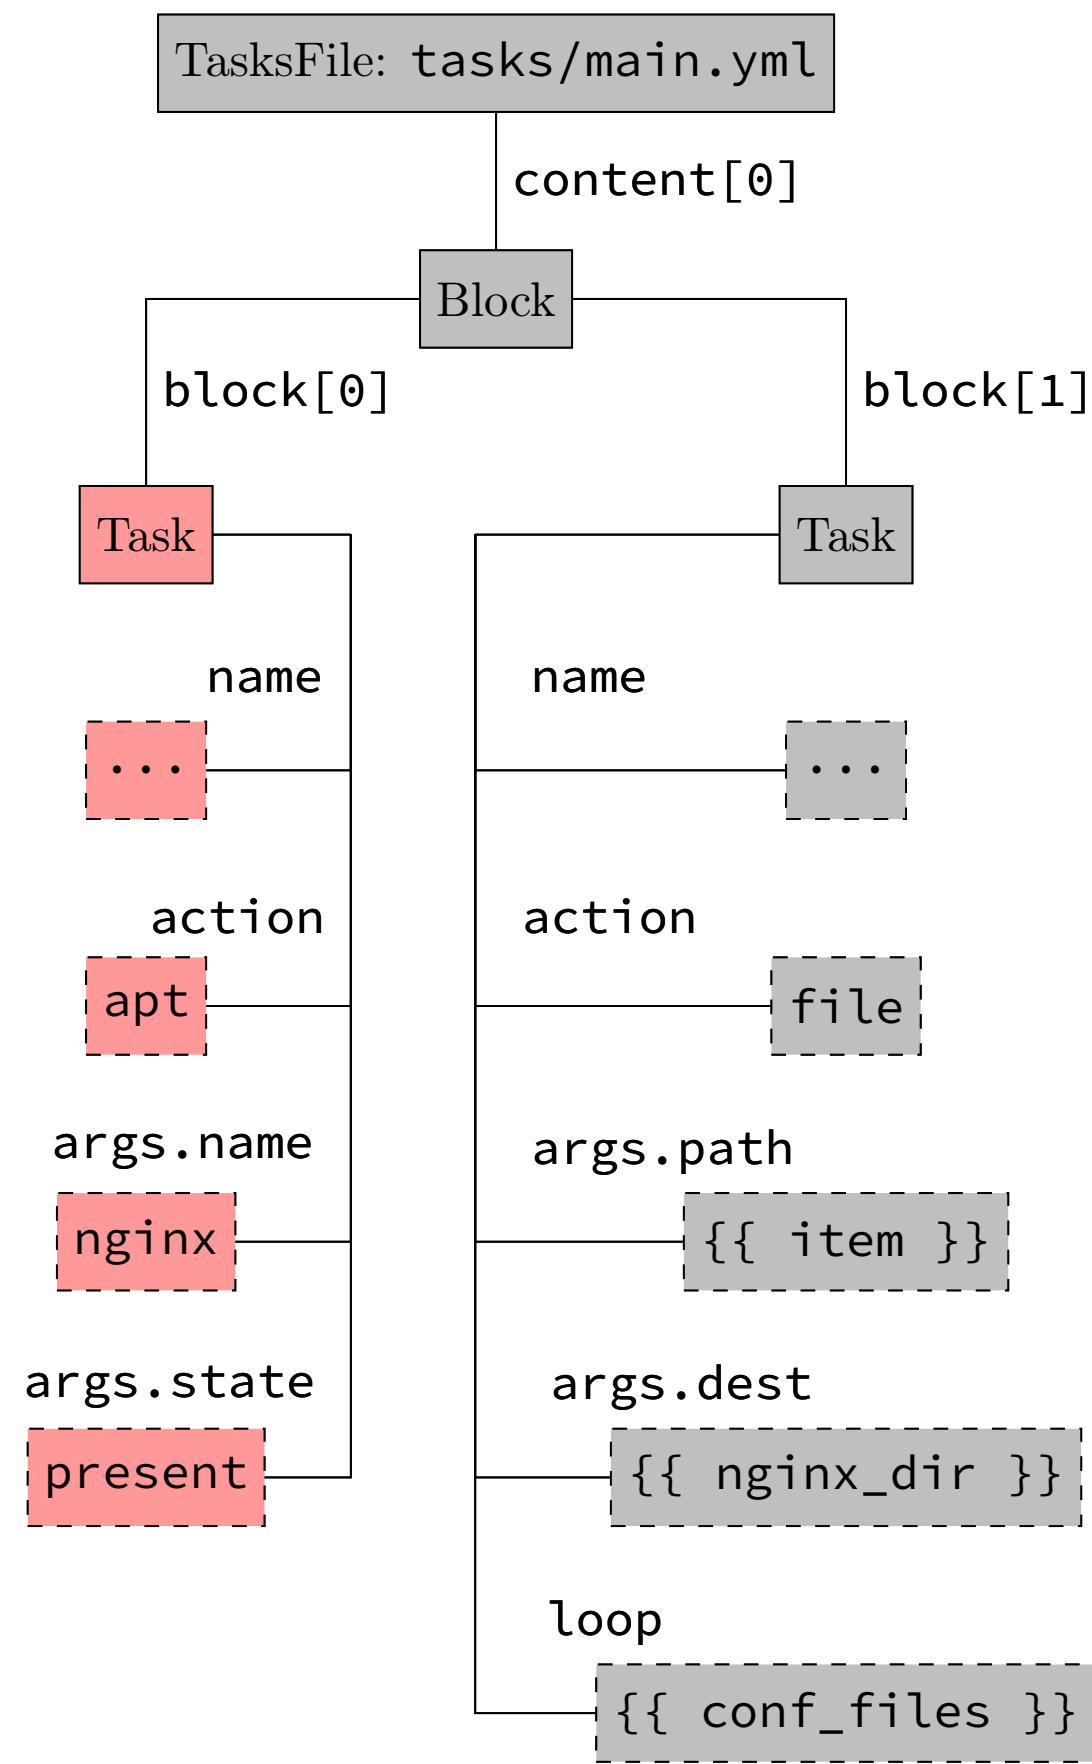


v1 v2

sim_{task} = 1


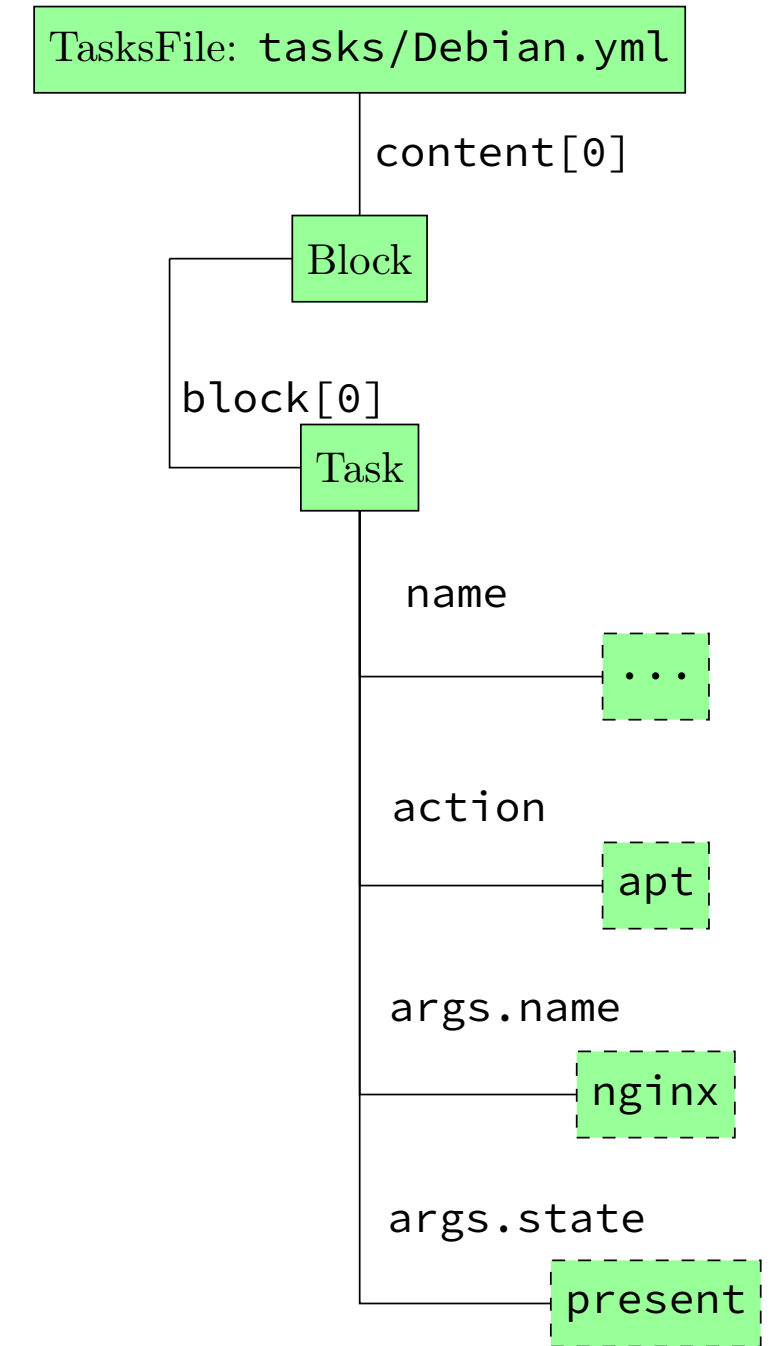
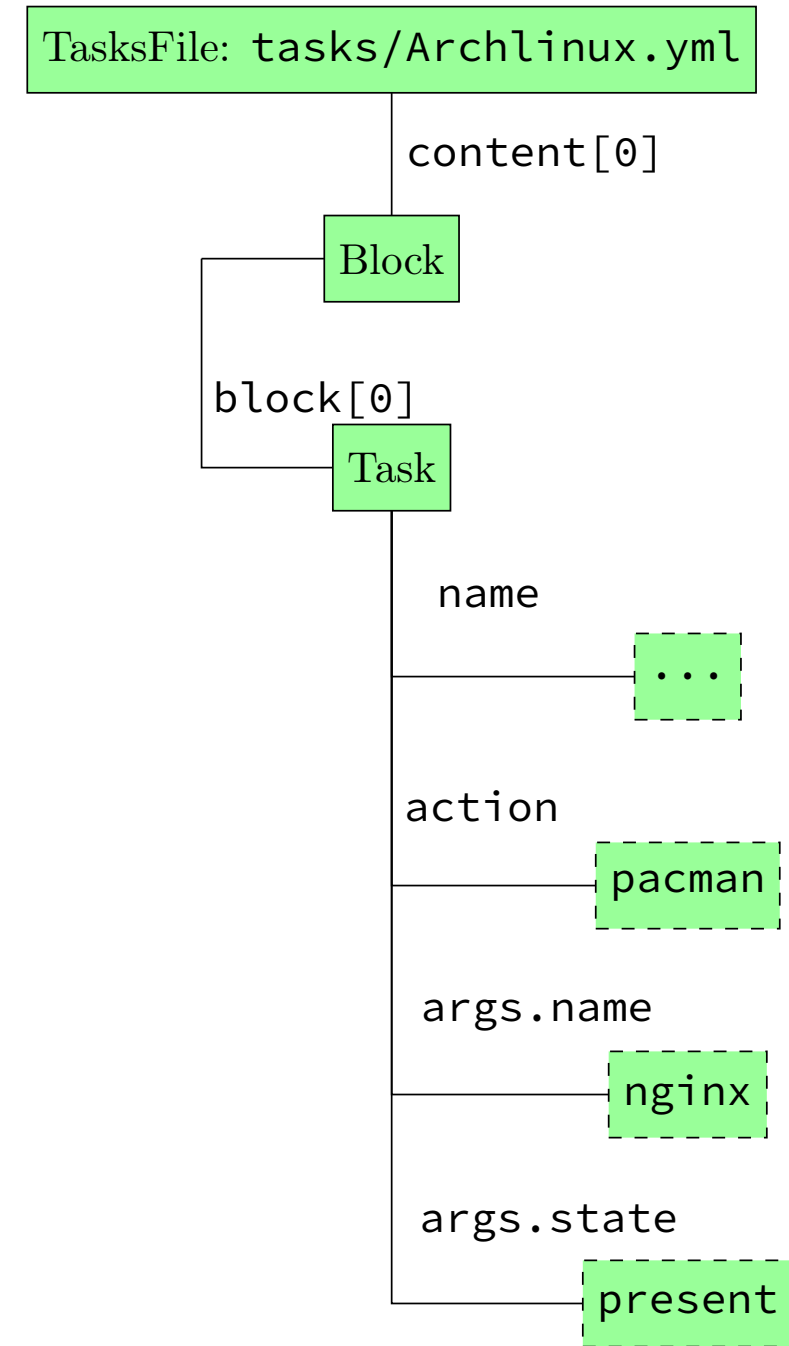



Task Matching



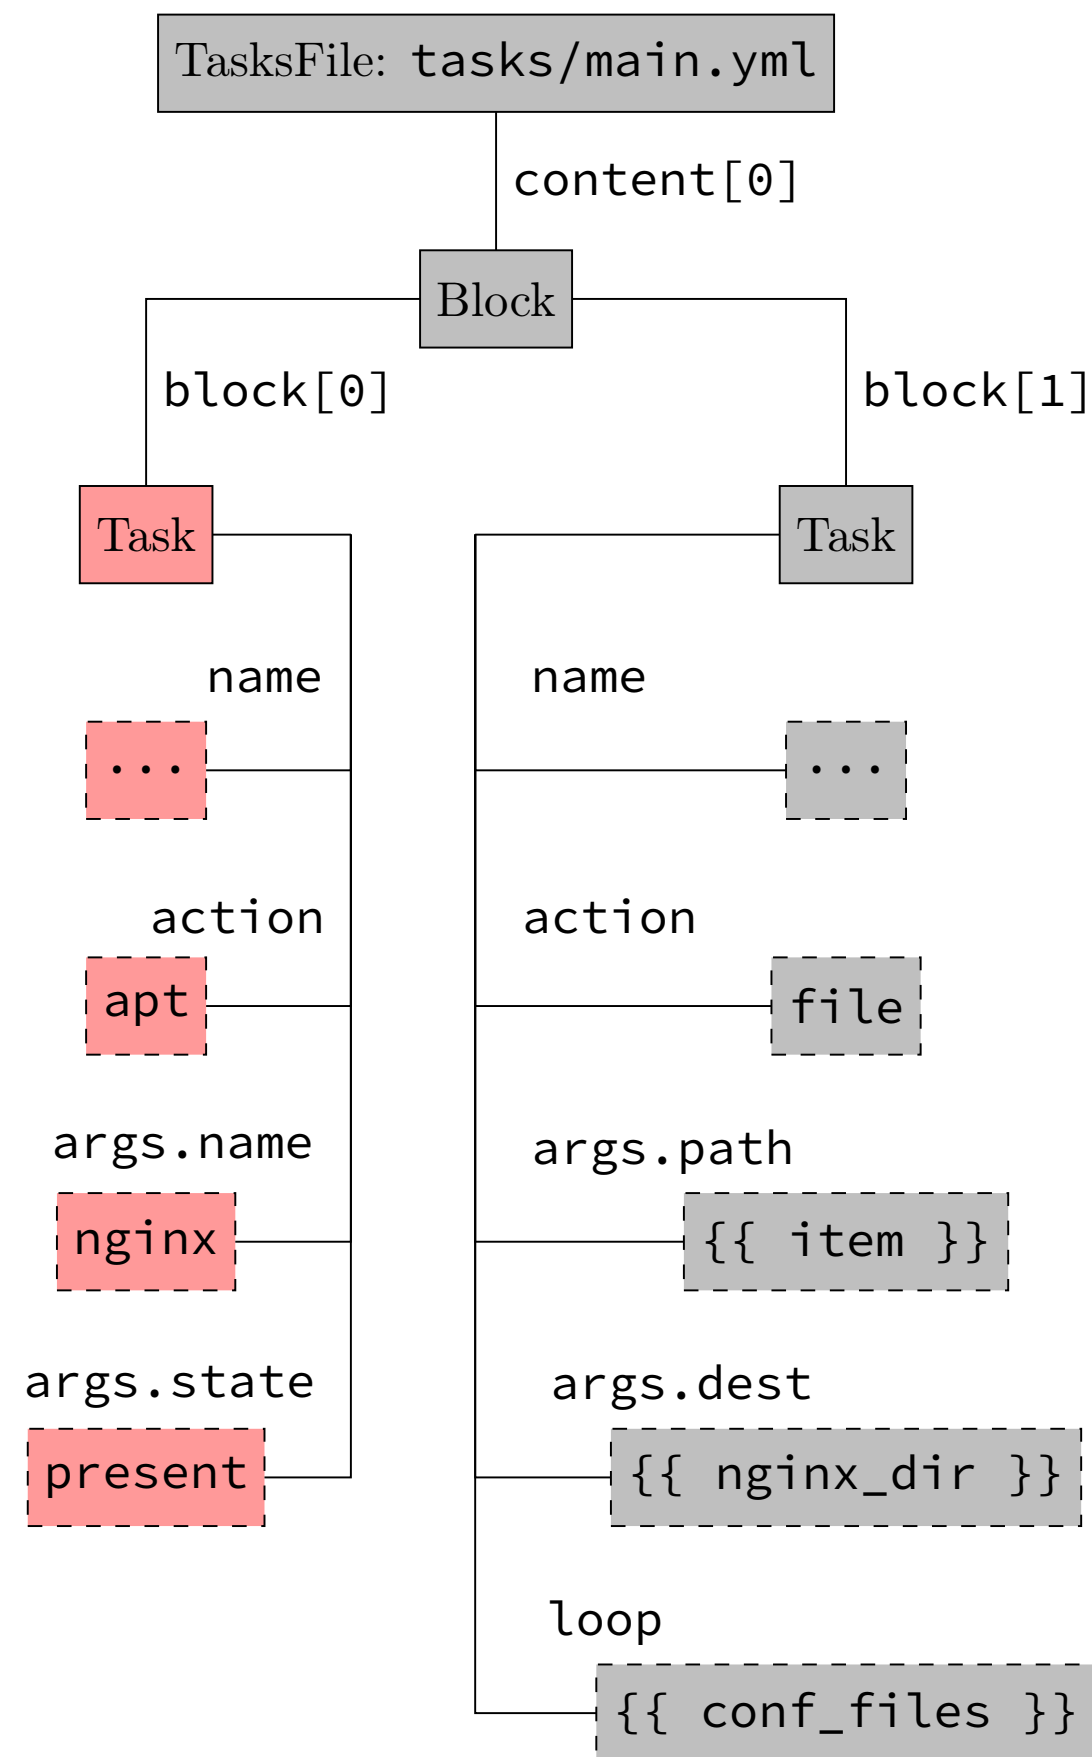
v1 v2

Task Matching

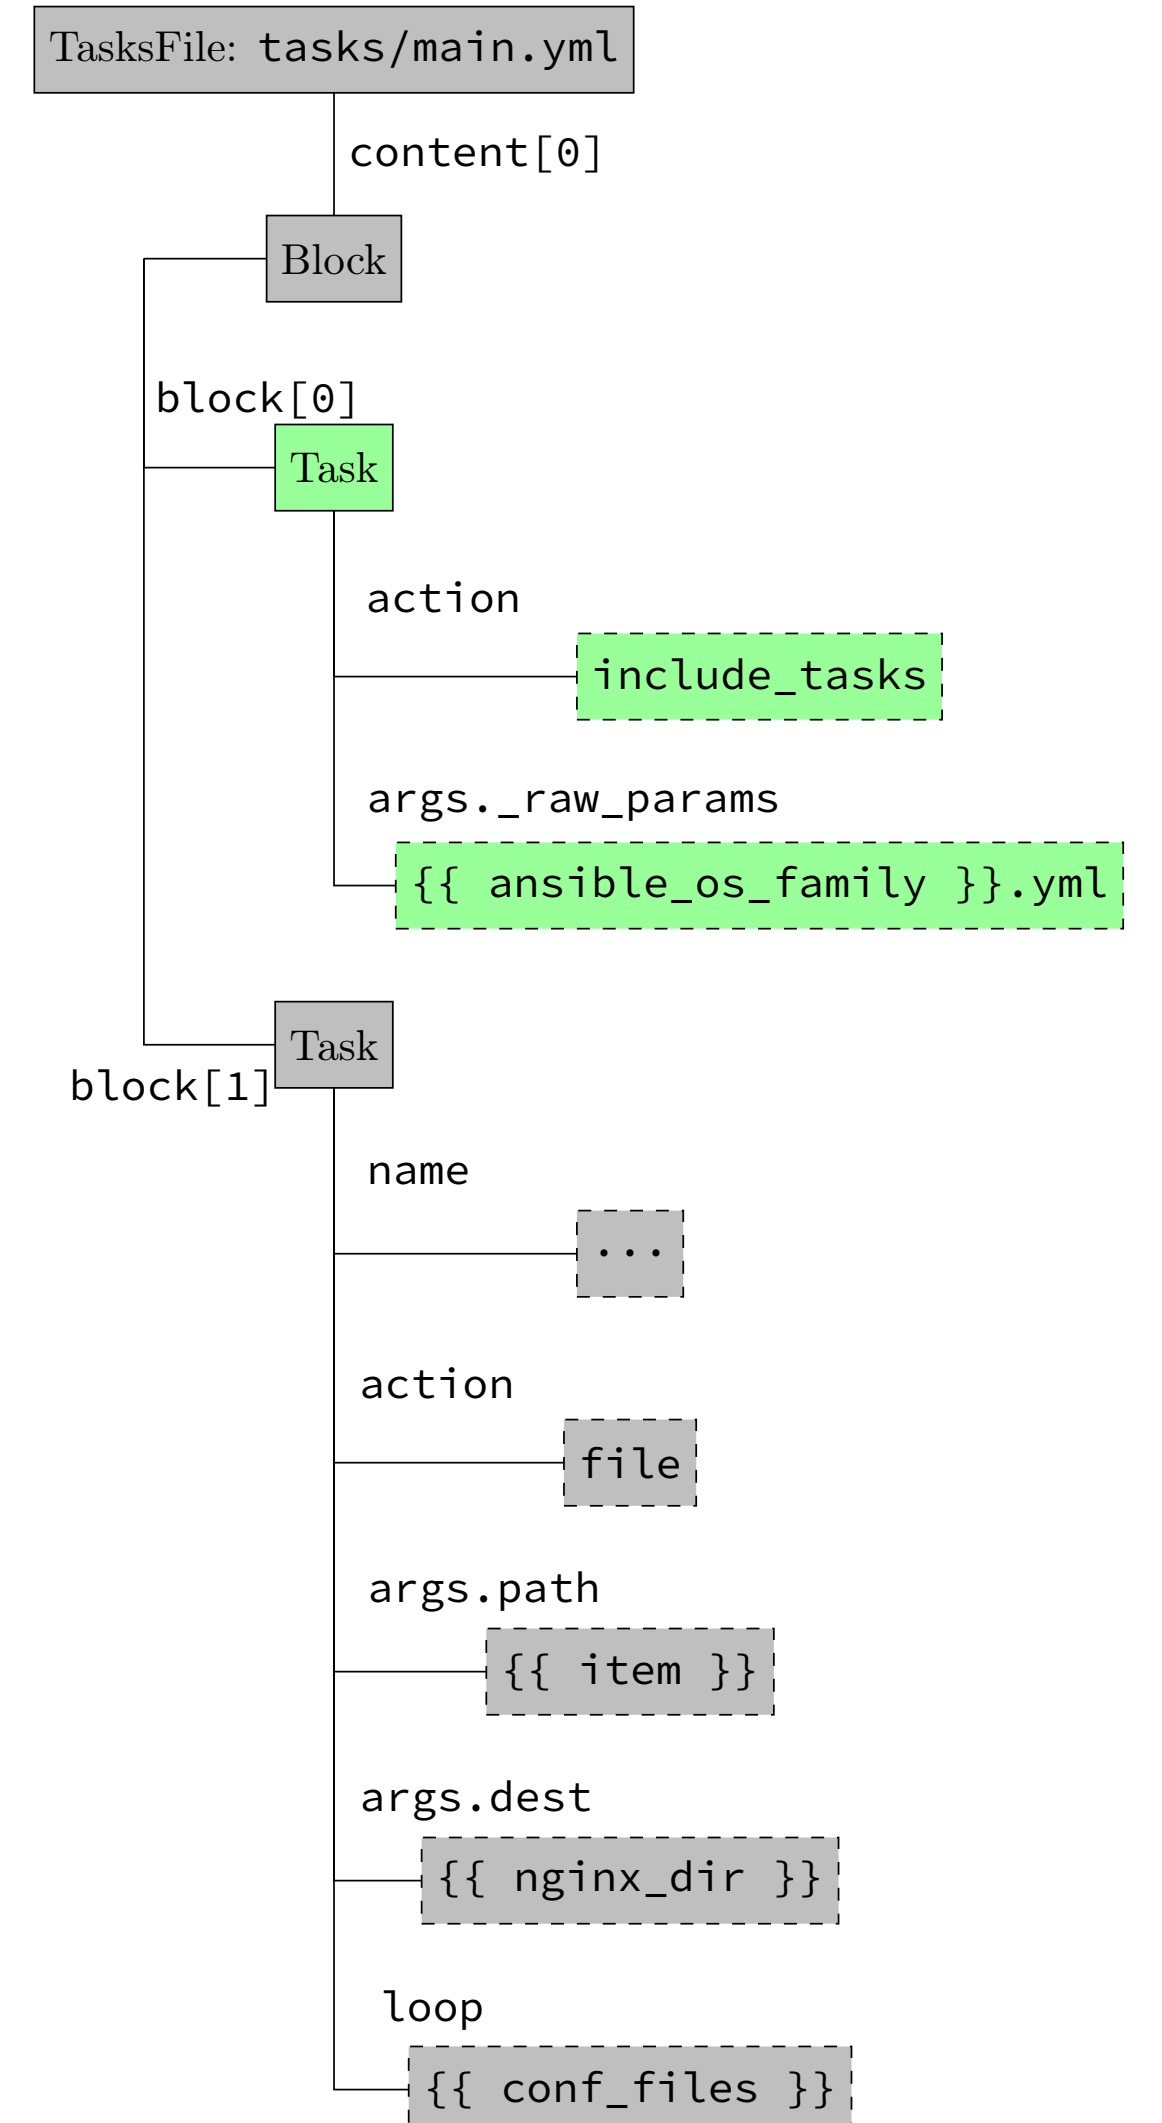
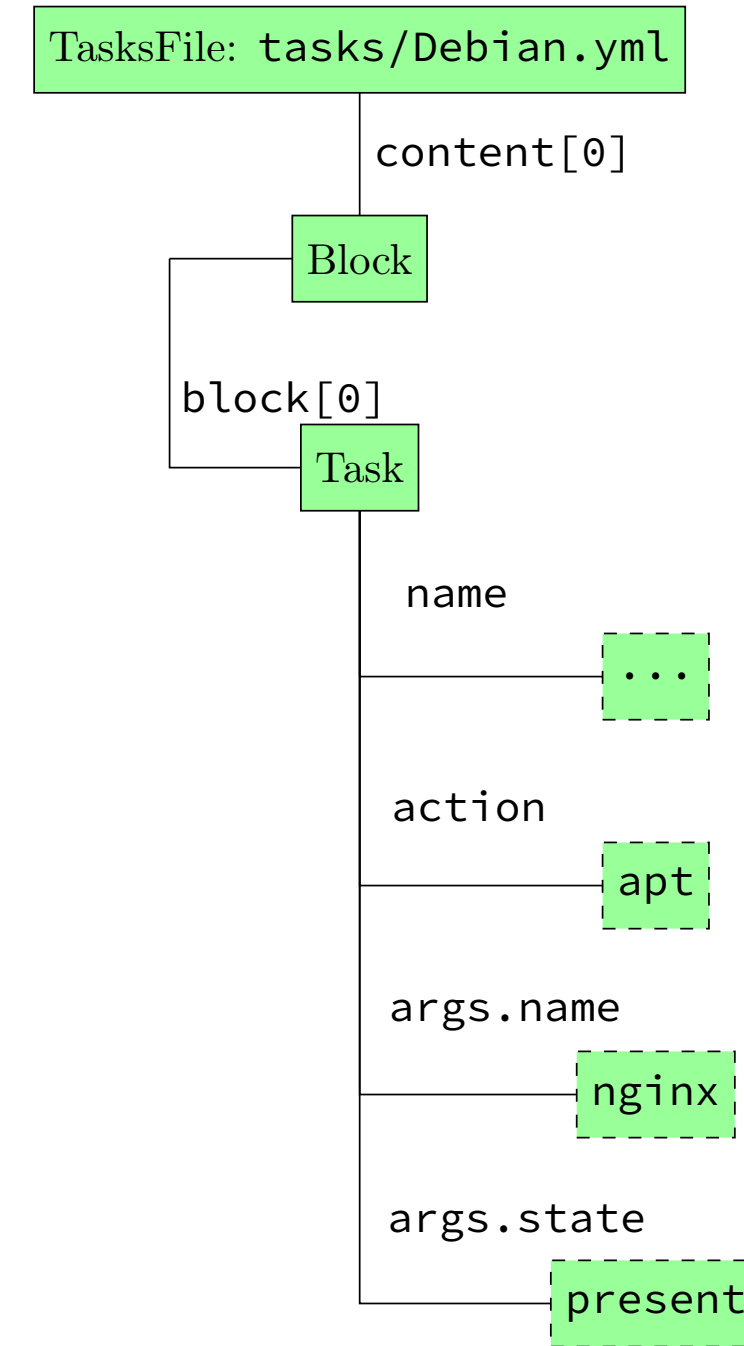
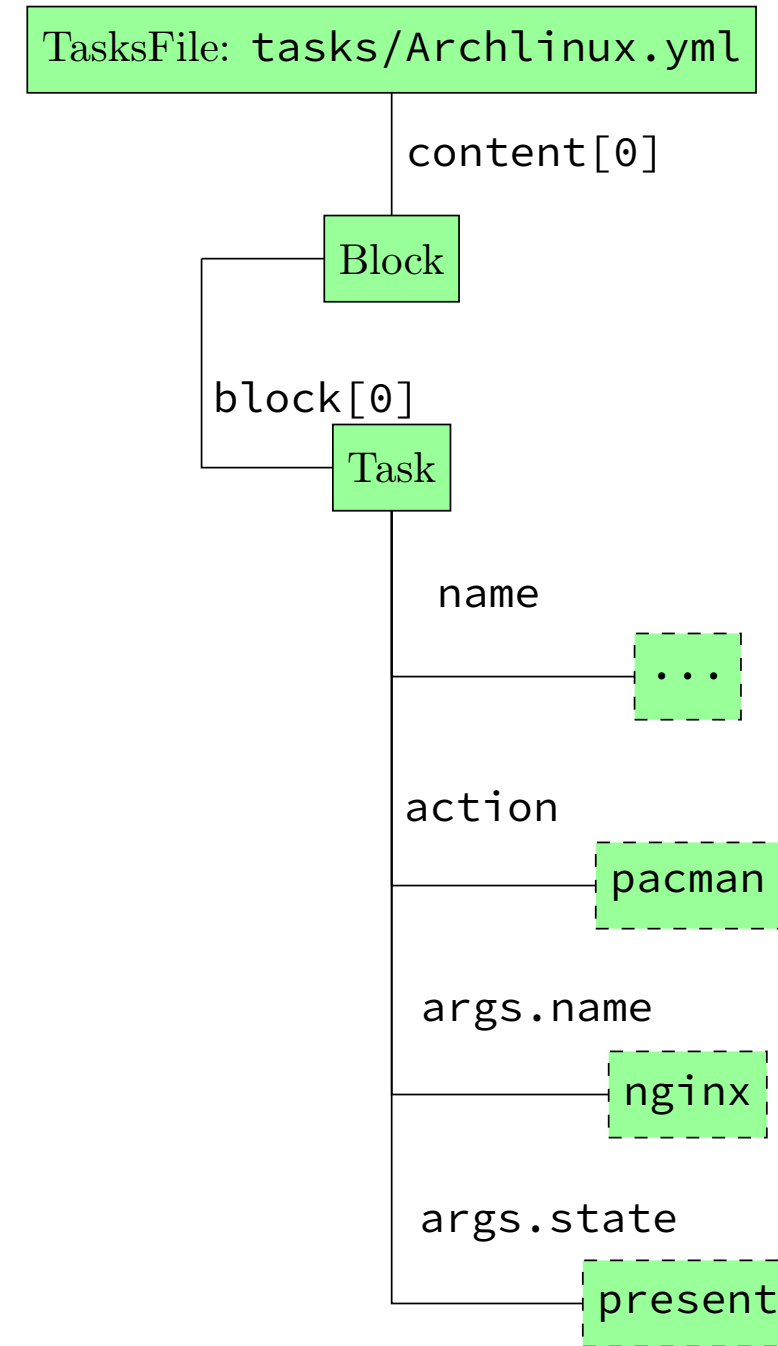


v1 v2

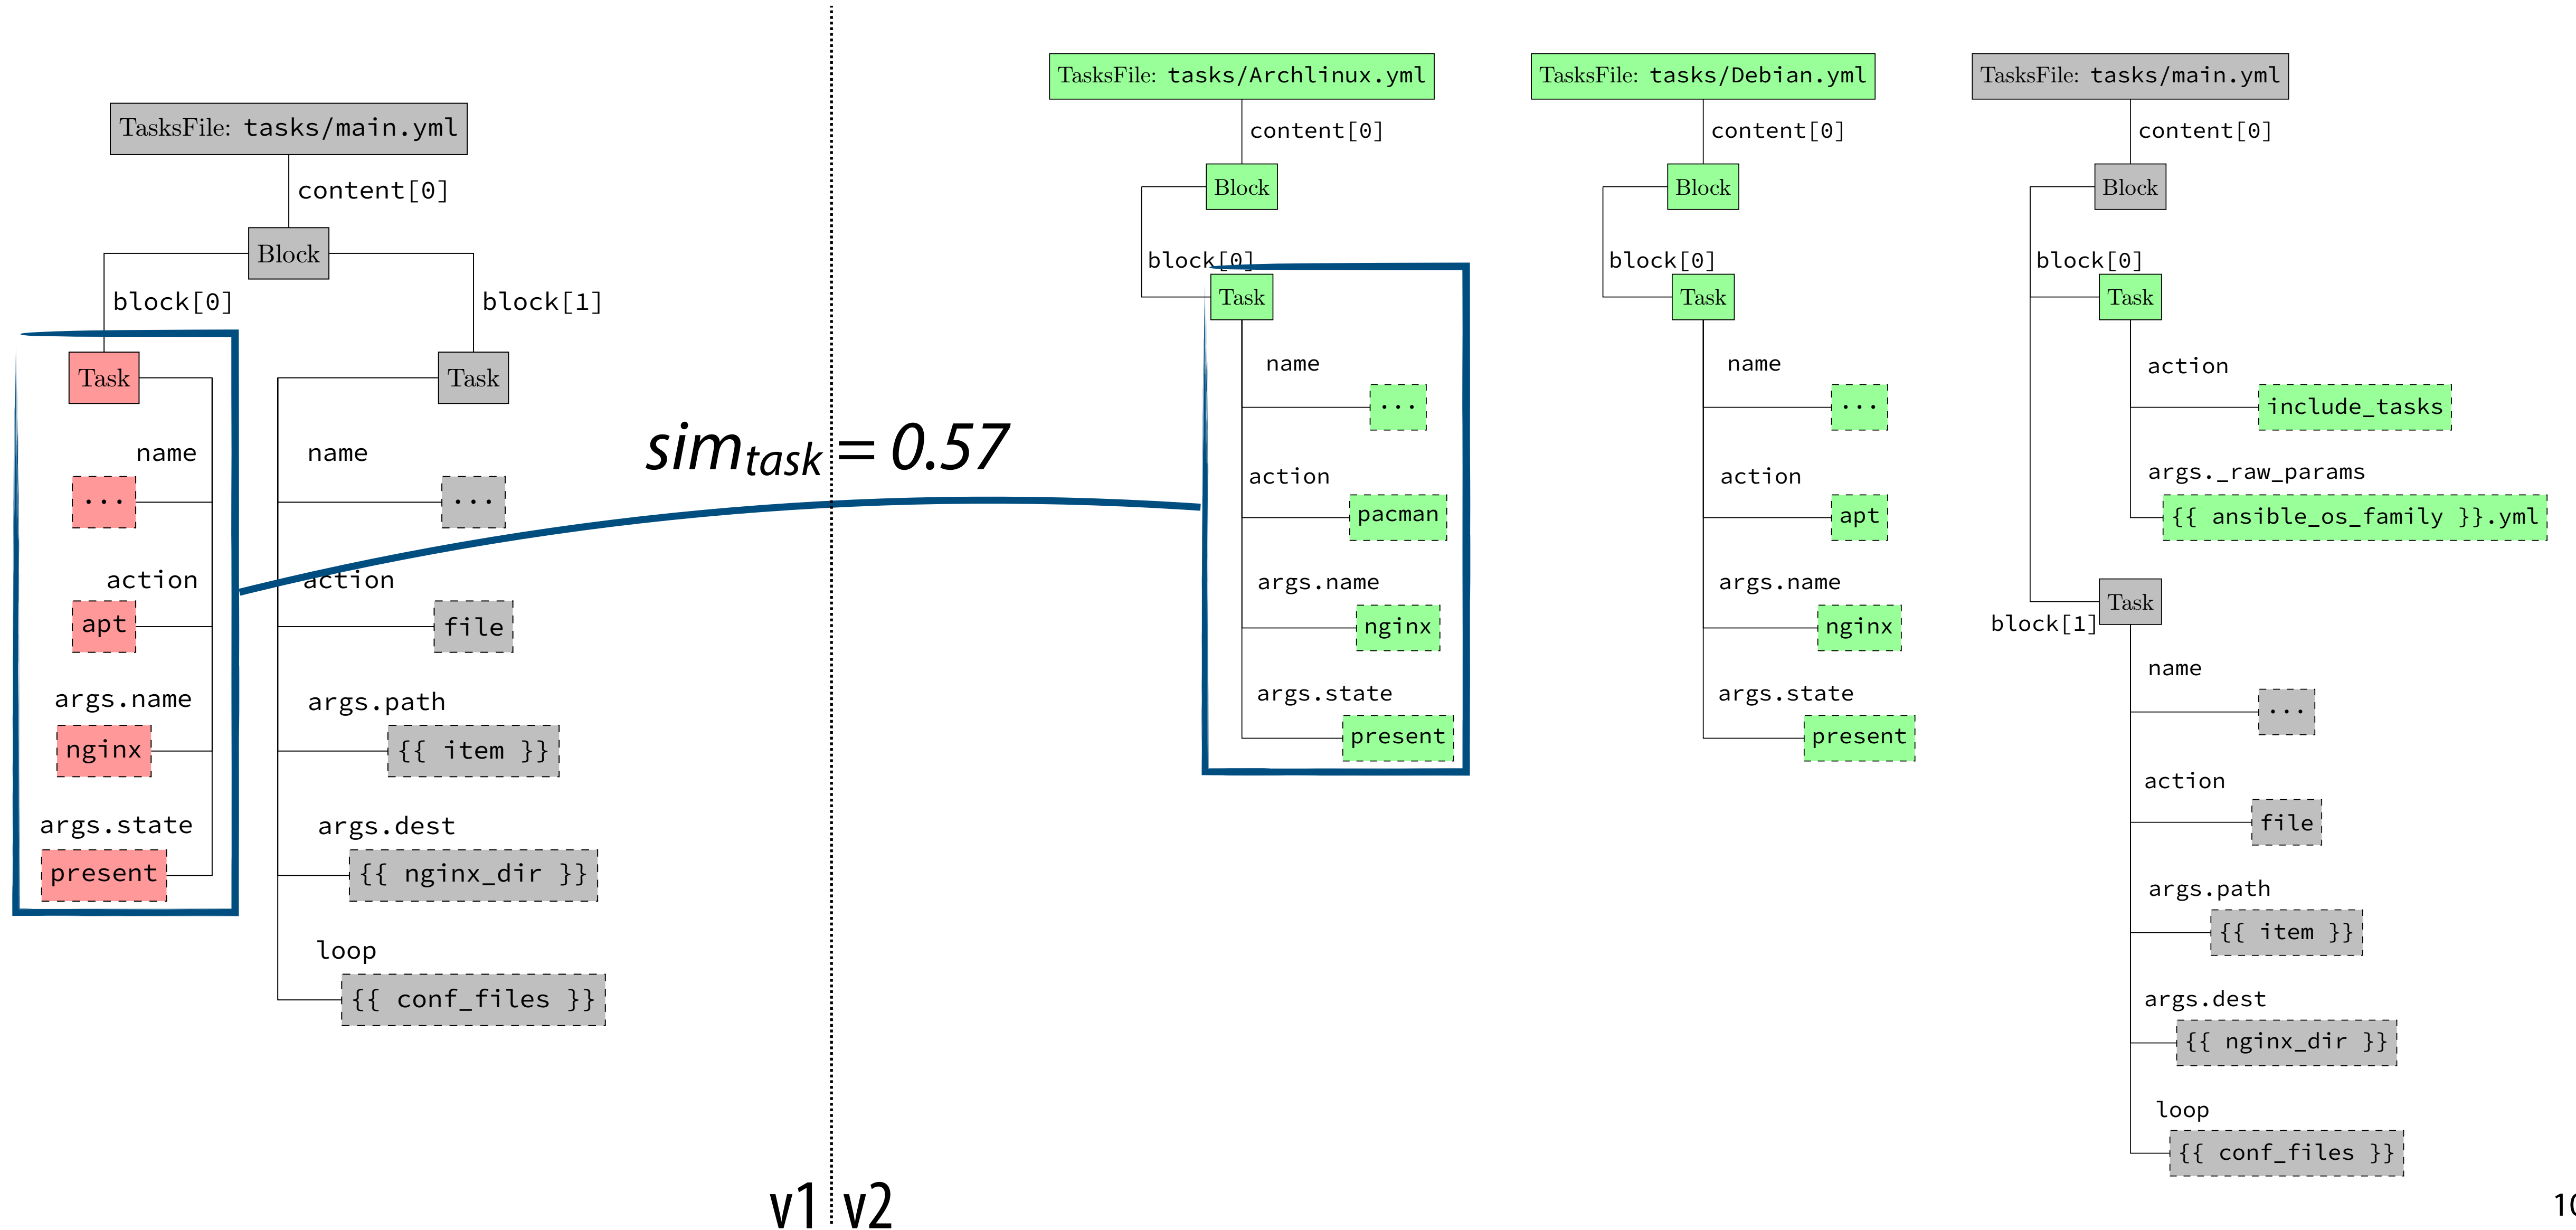
Task Relocation Matching



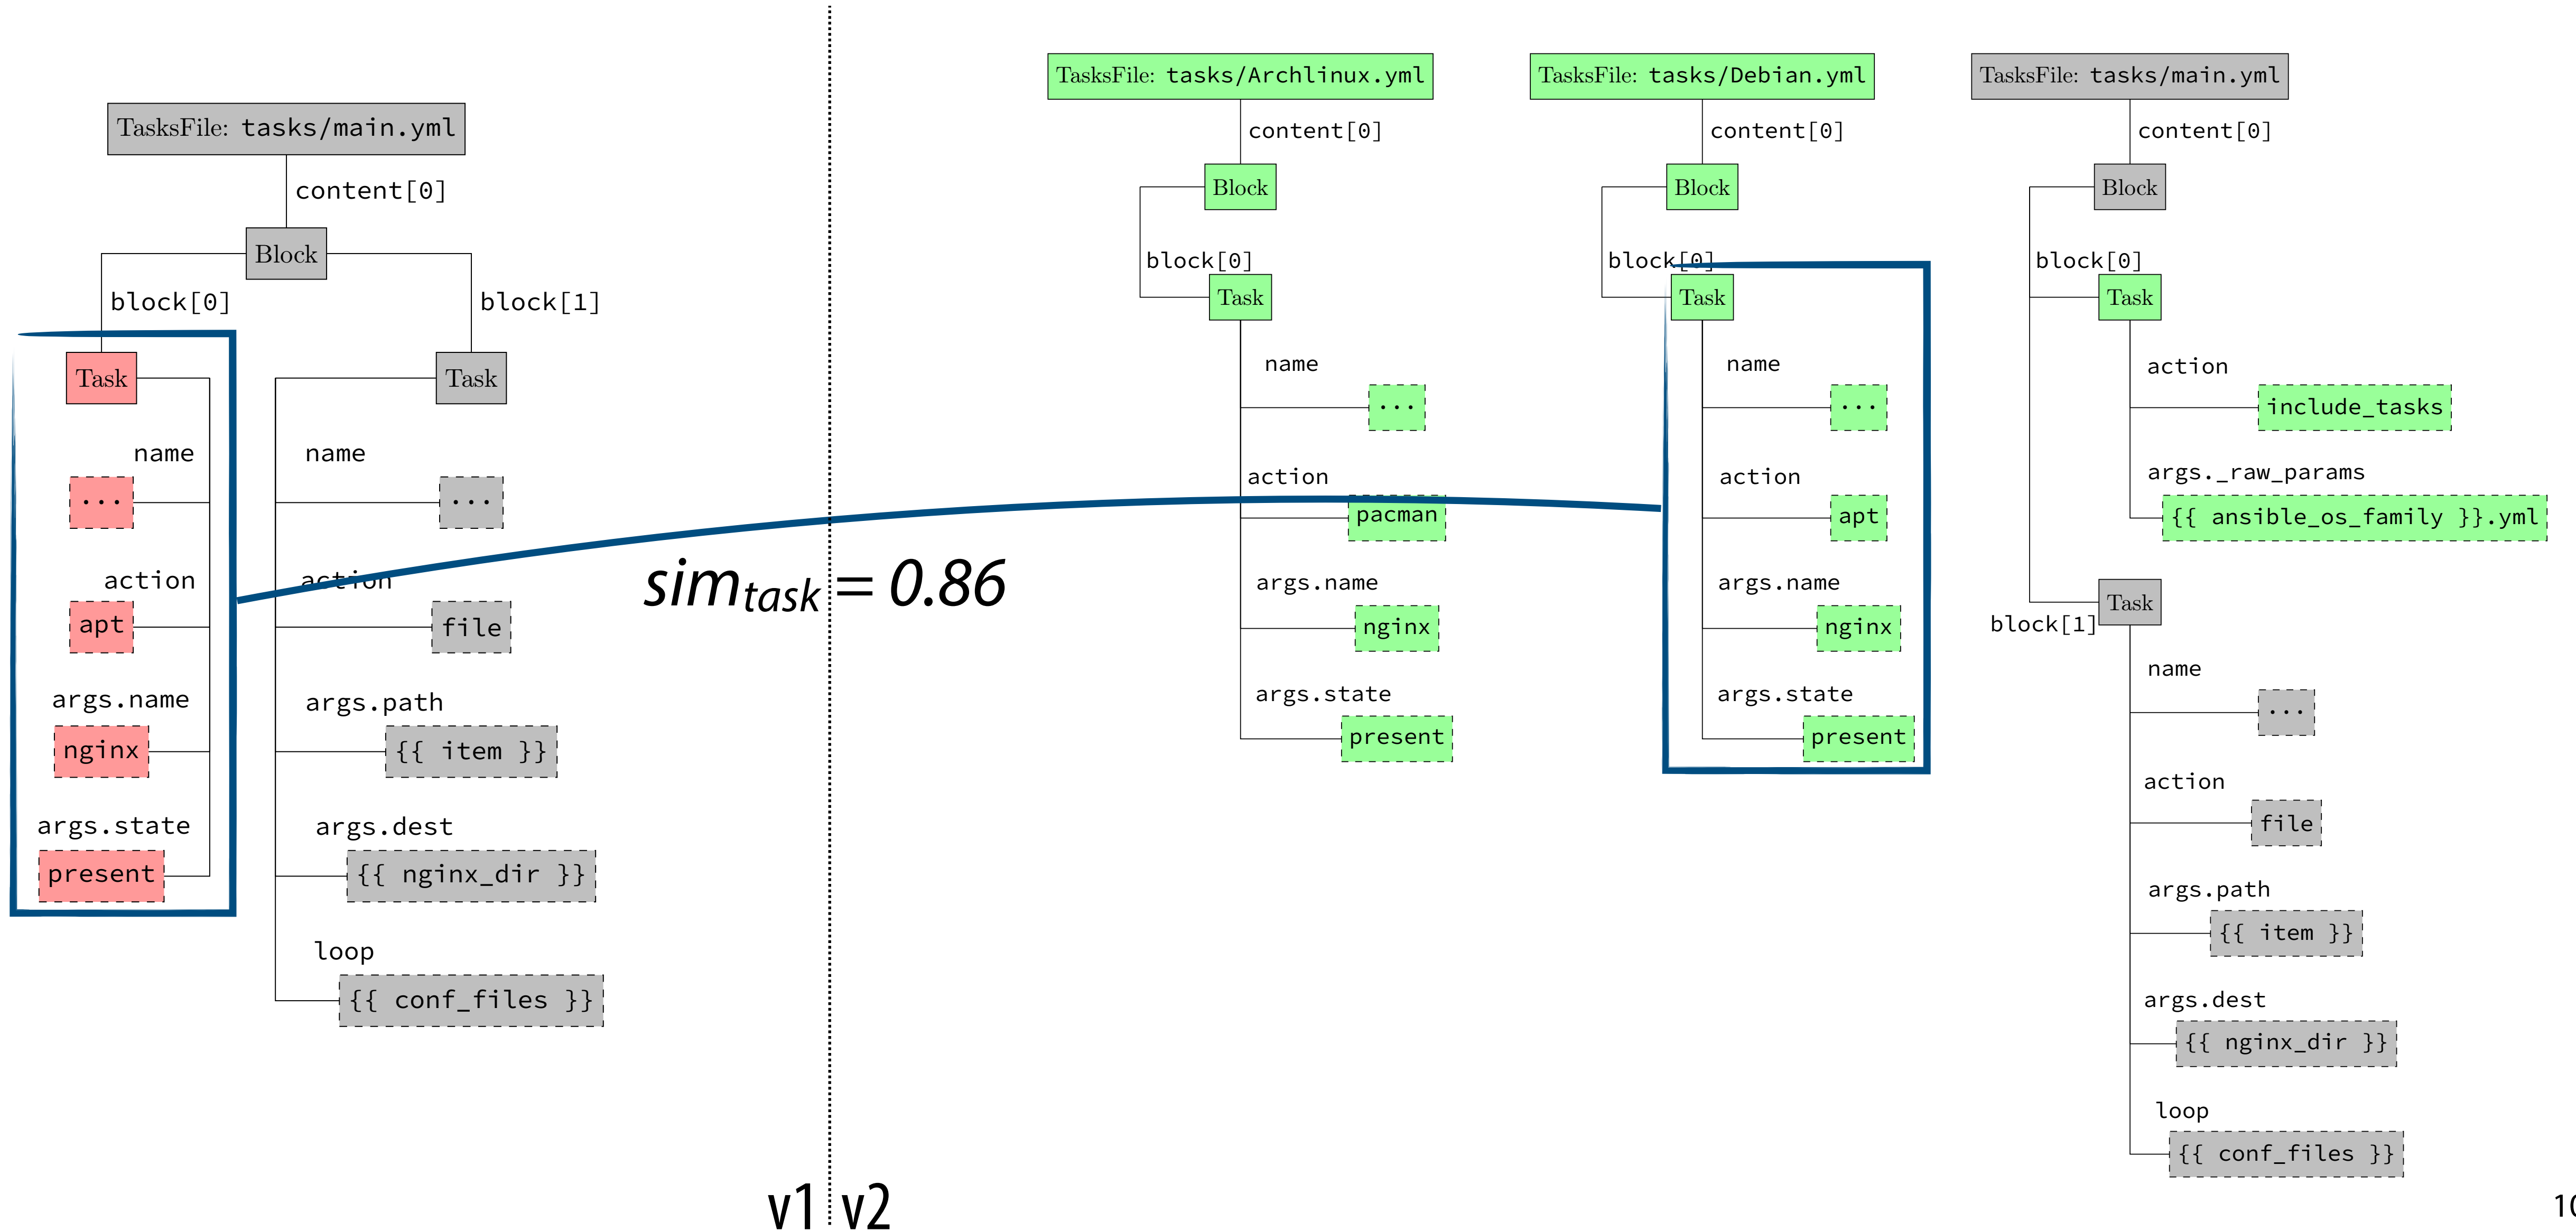
v1 v2



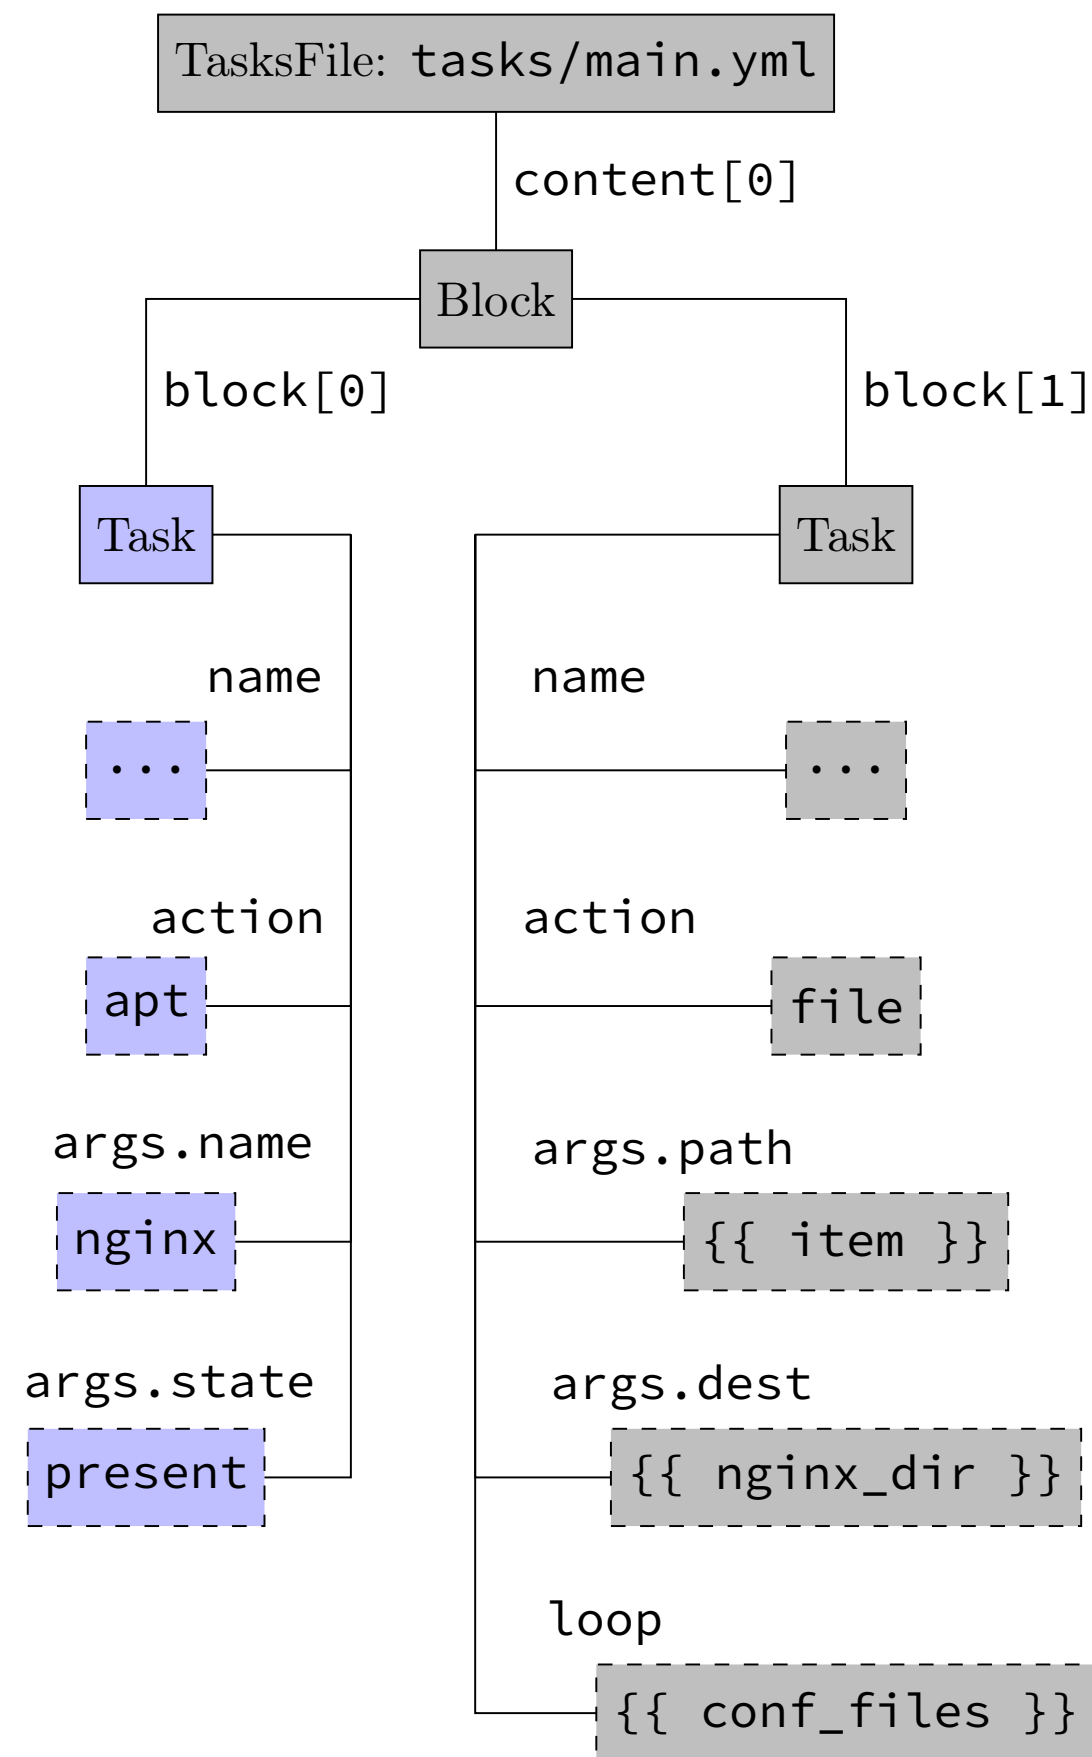
Task Relocation Matching



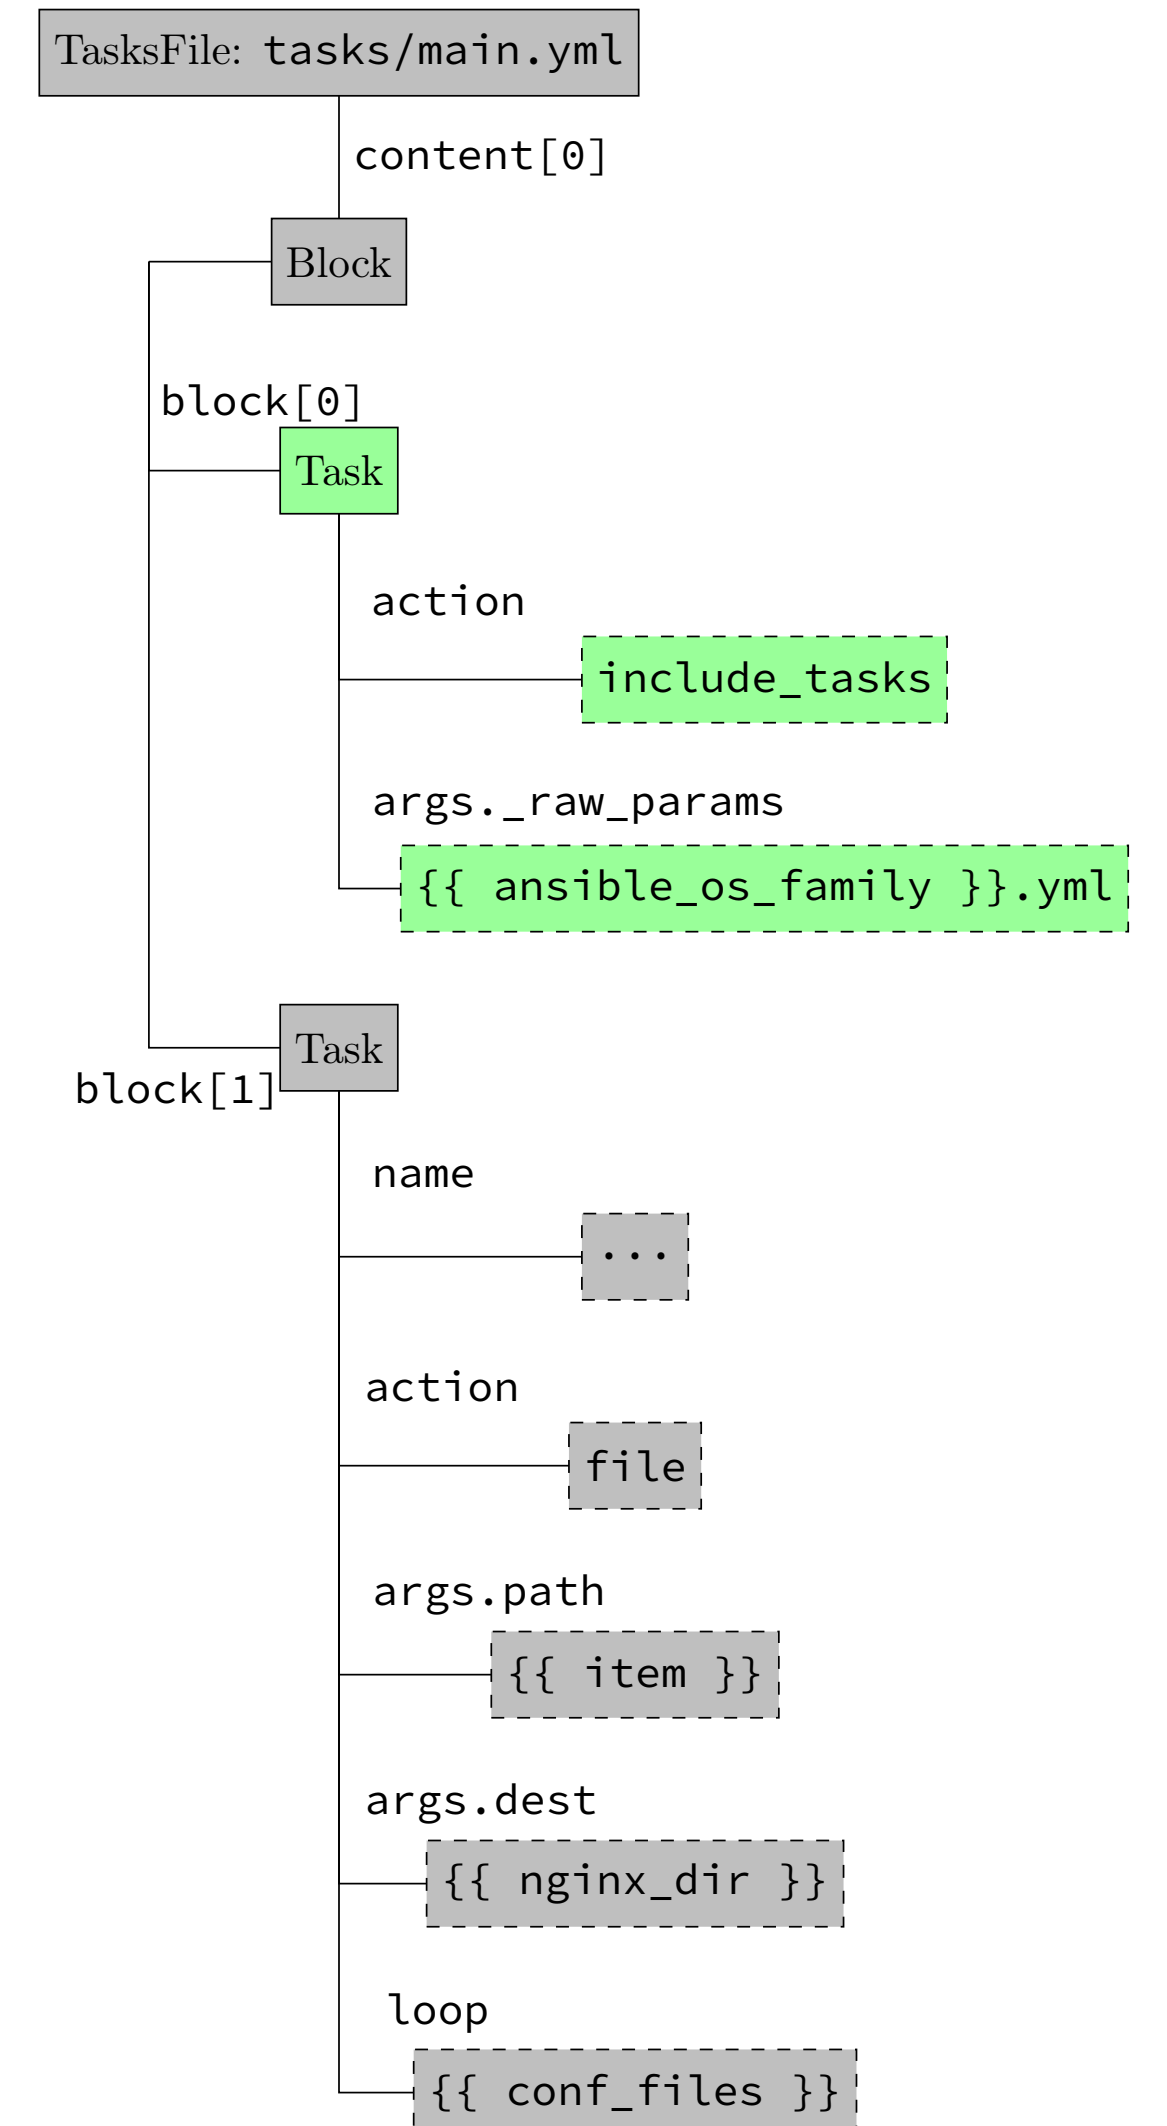
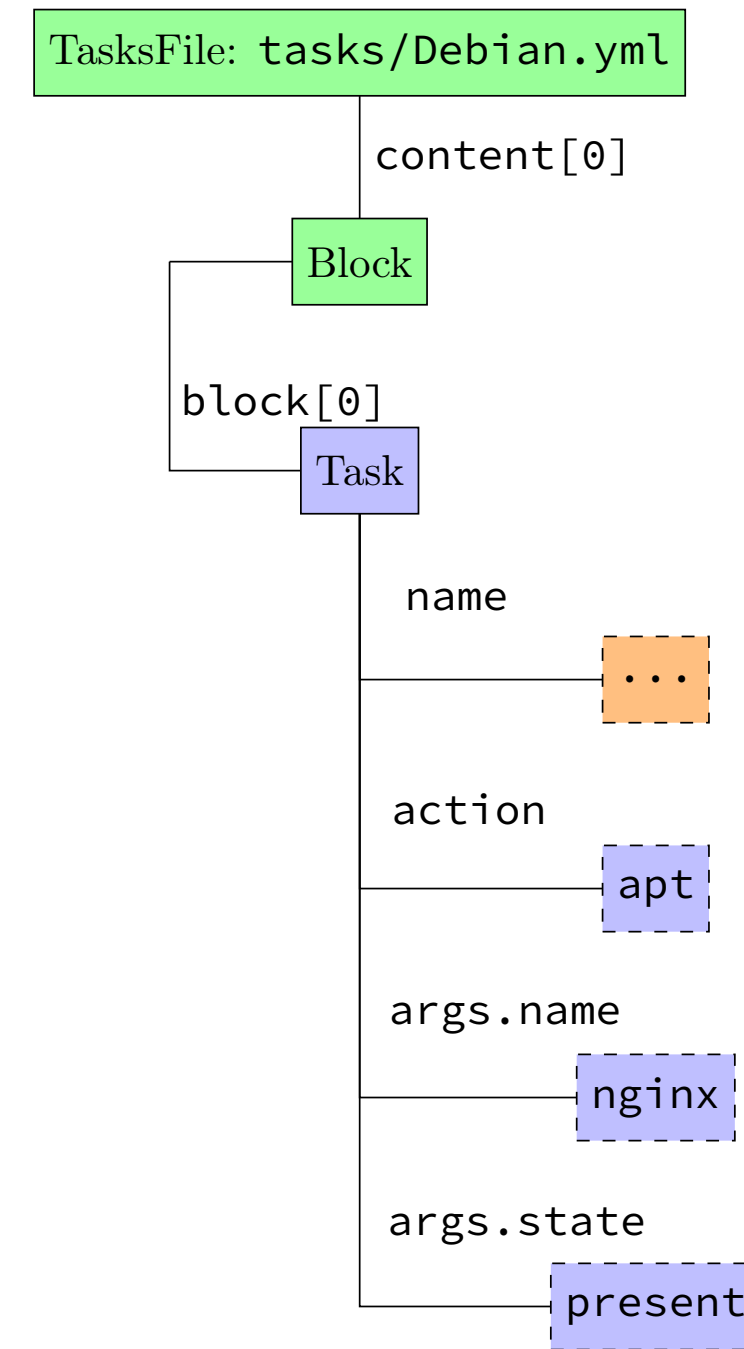
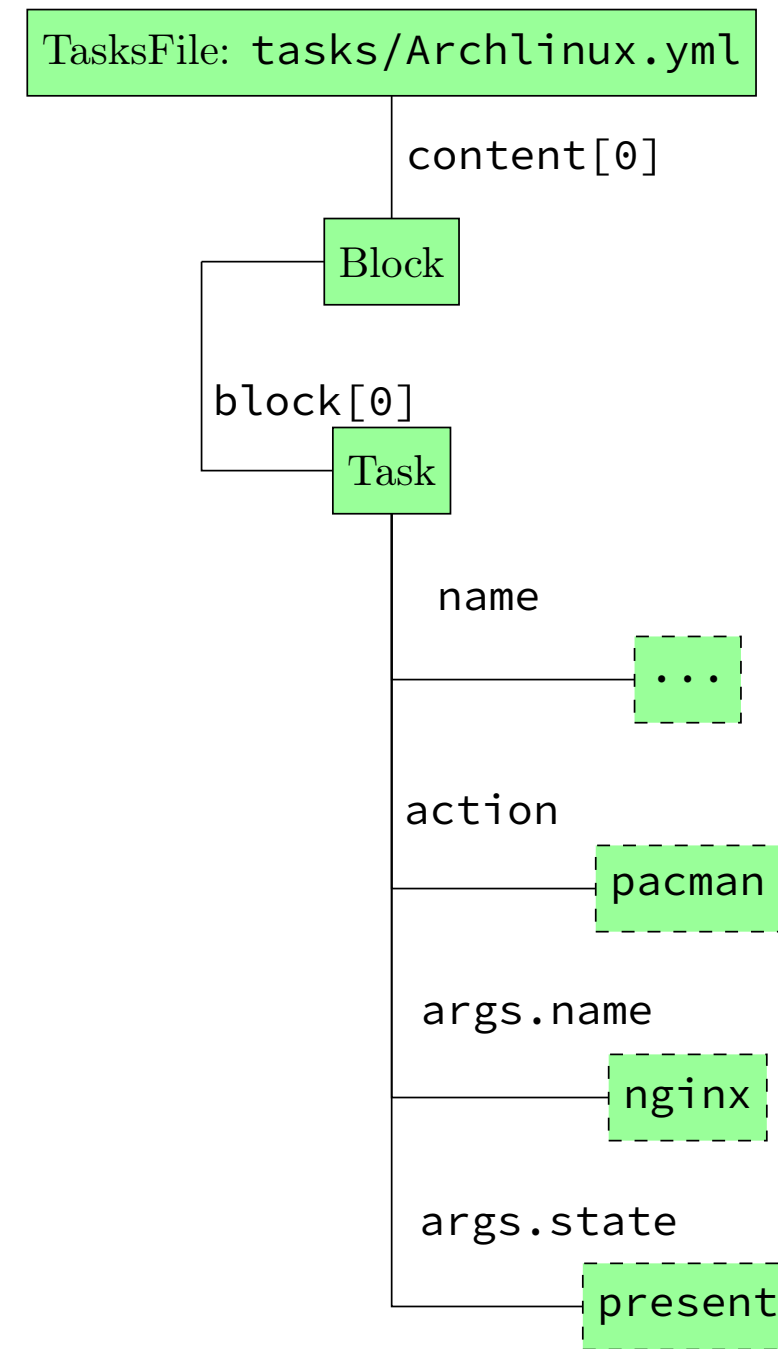
Task Relocation Matching



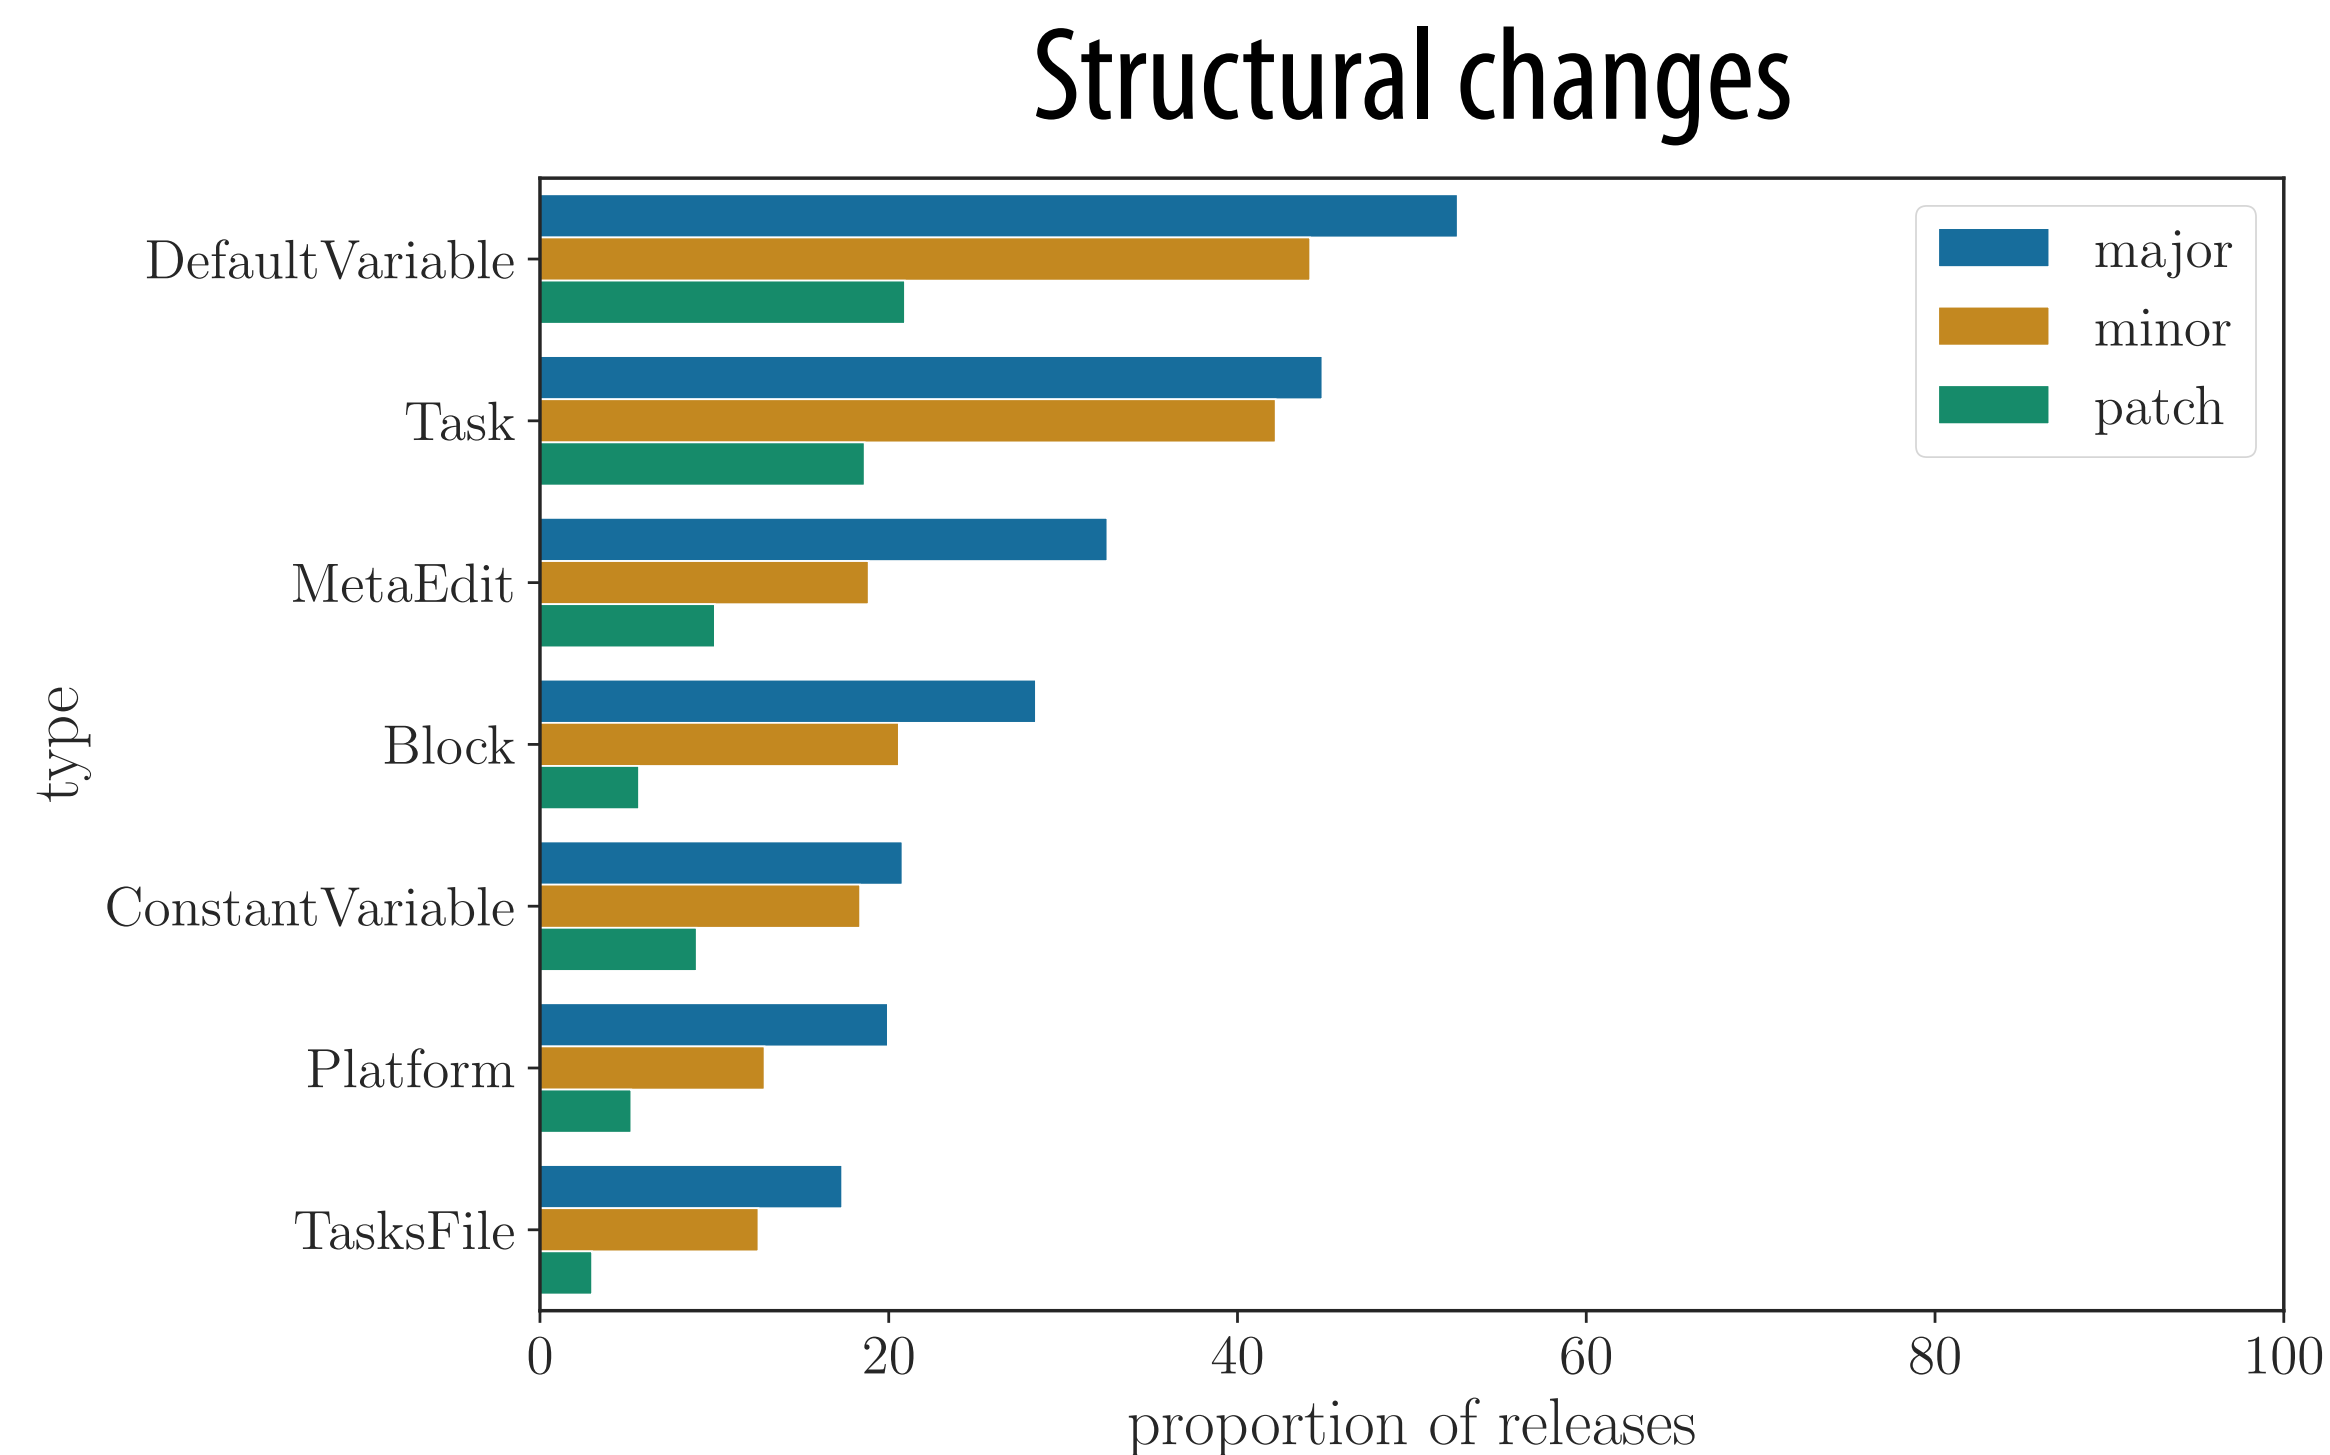
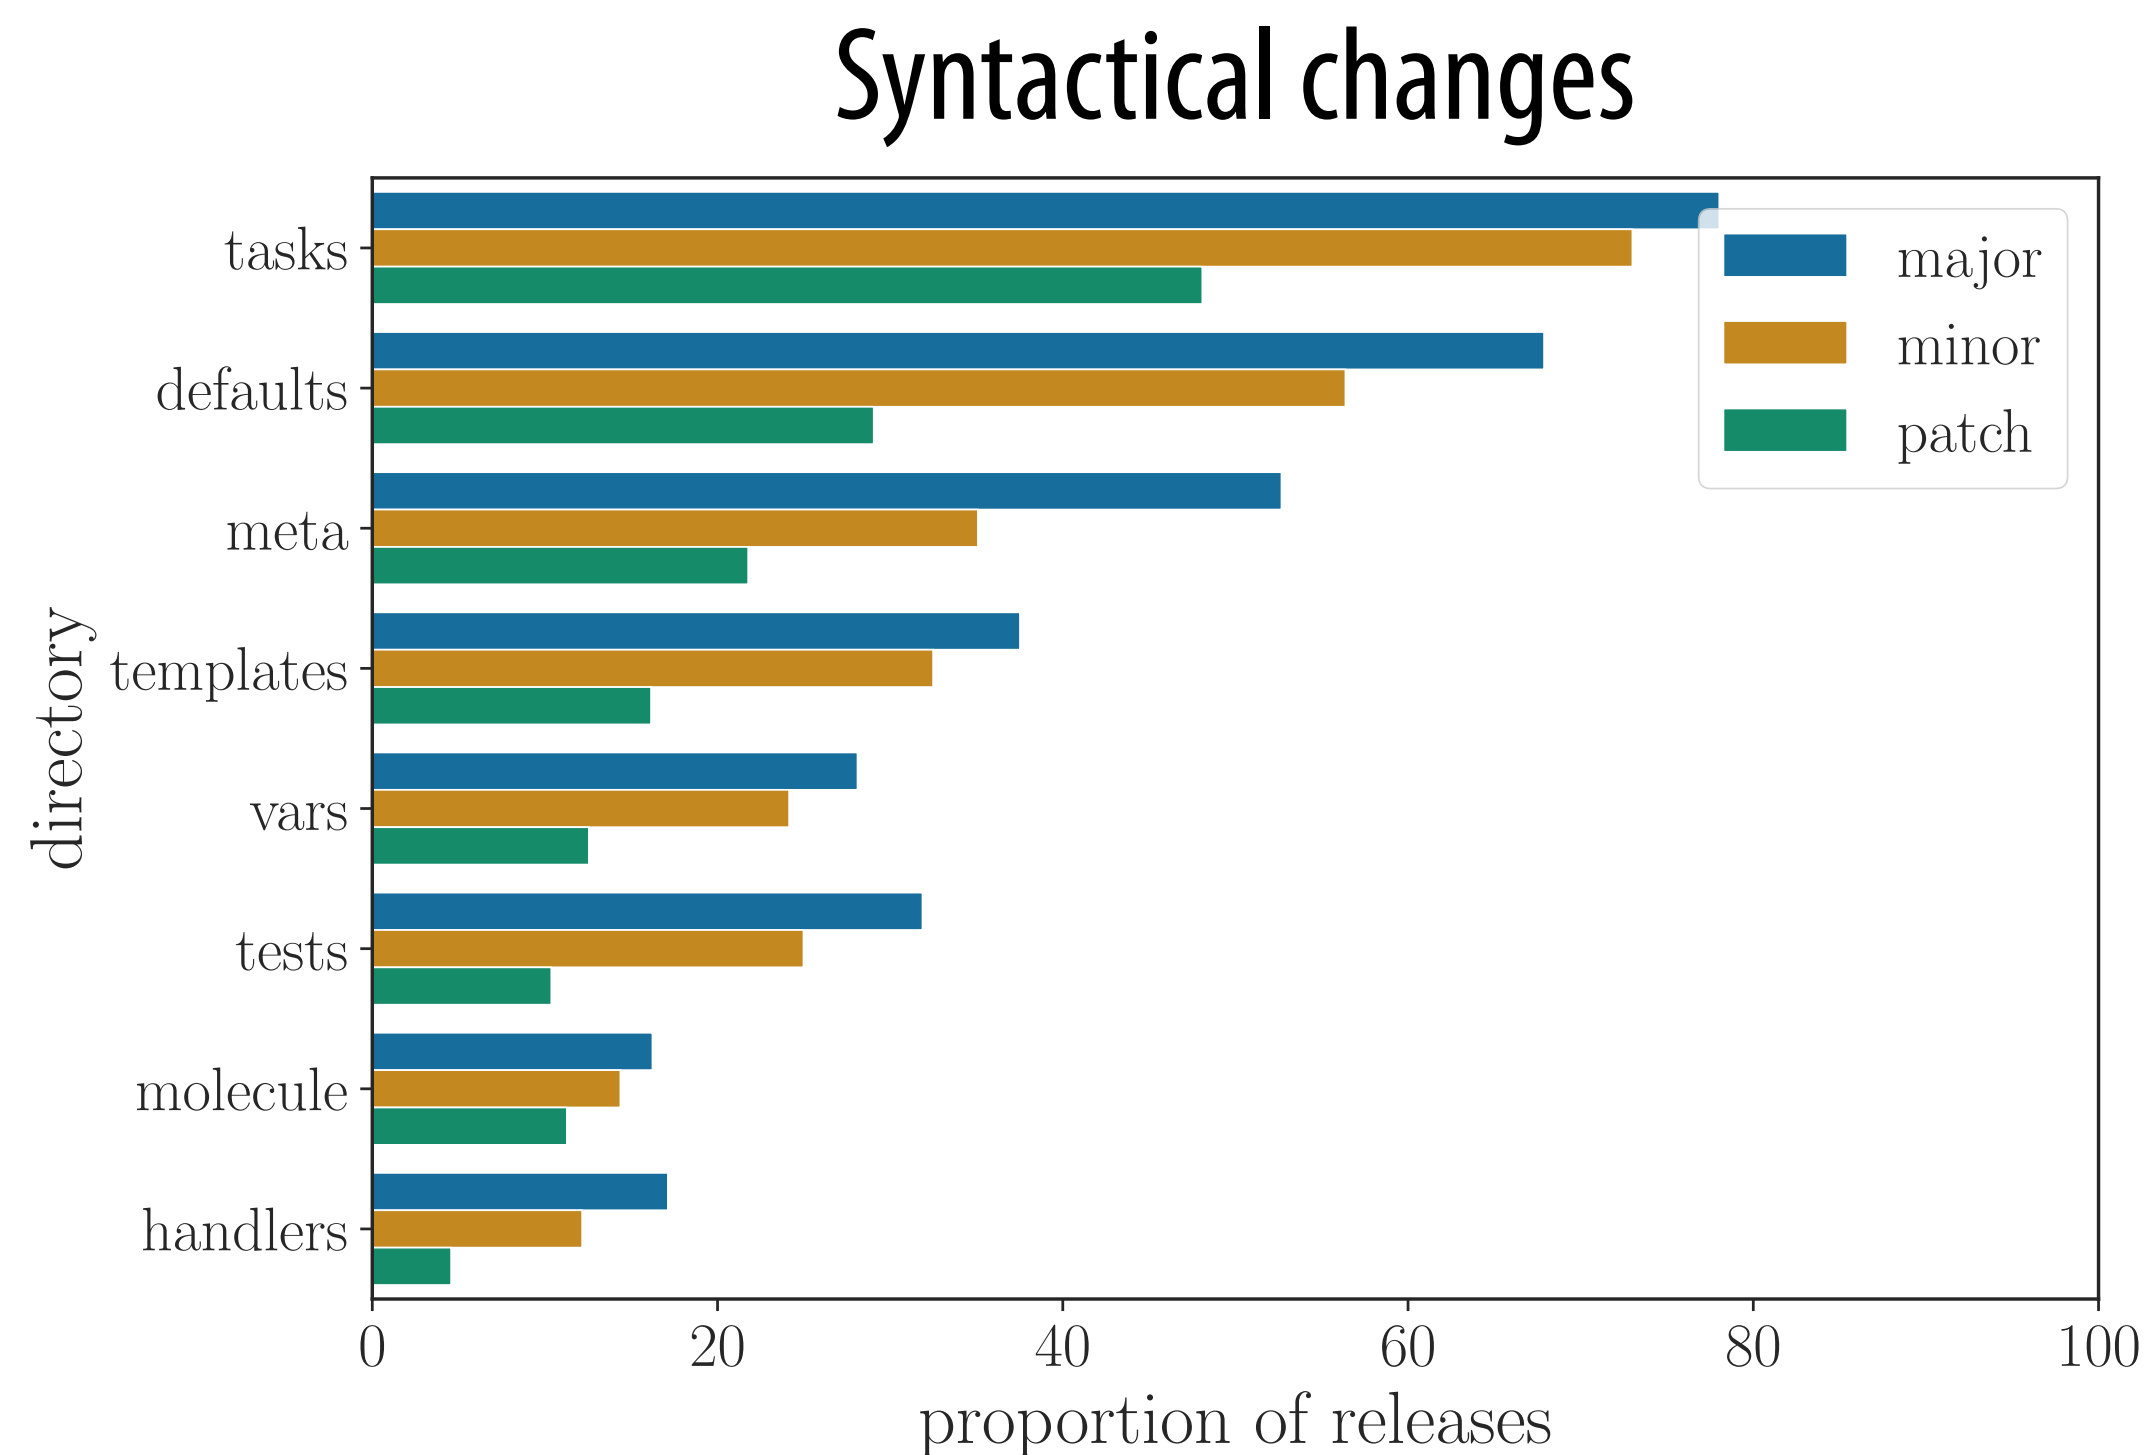
Task Relocation Matching



v1 v2



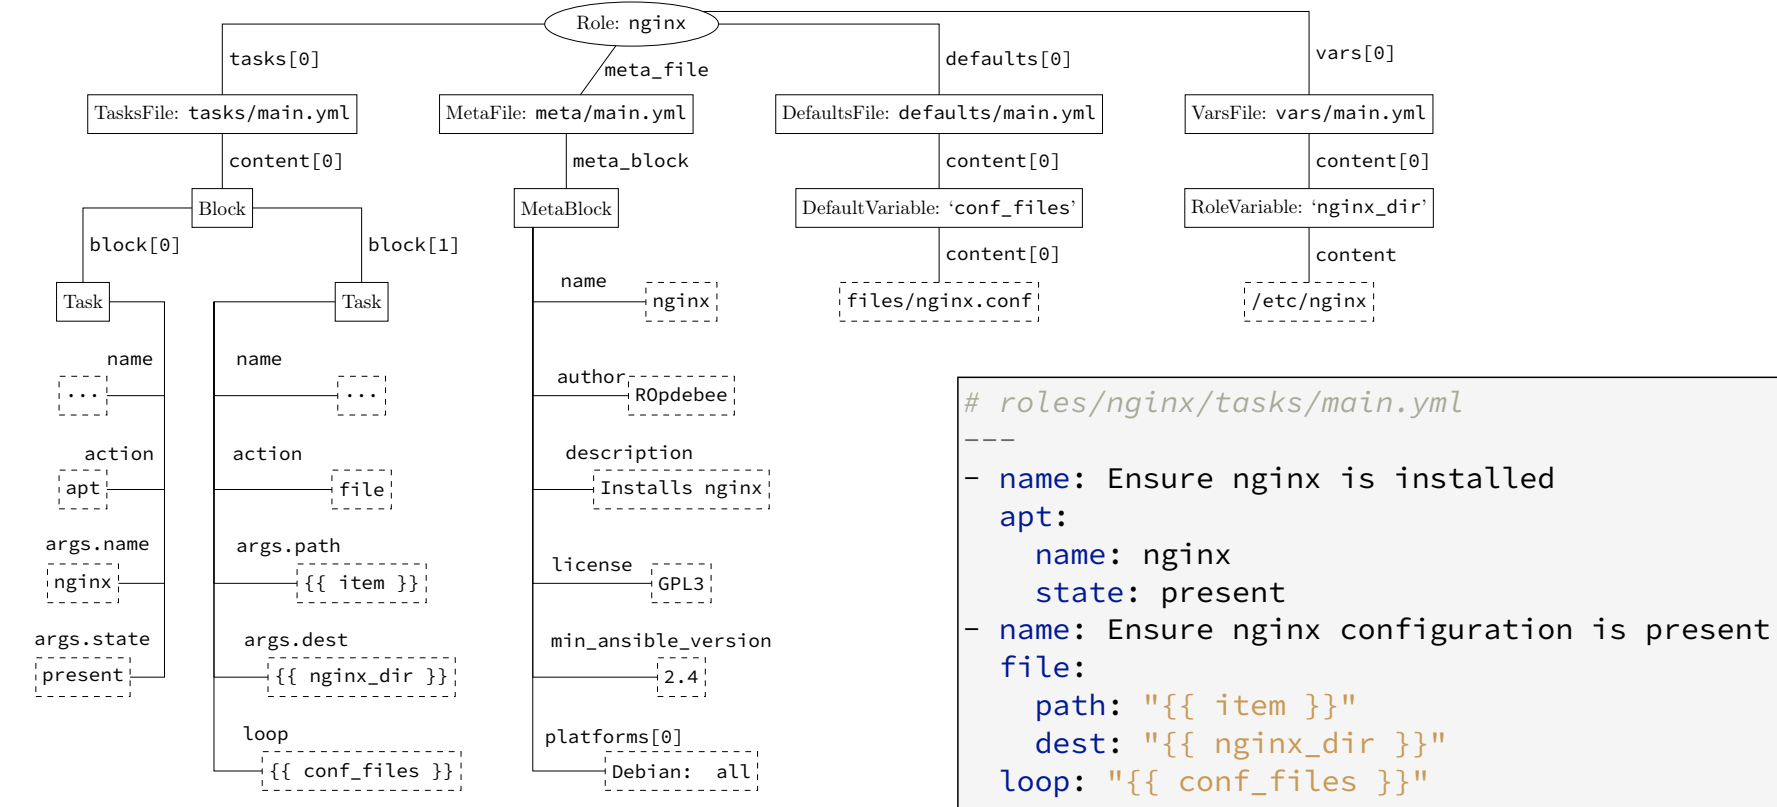
Structural Changes in Role Releases



~30% of releases contain no structural change

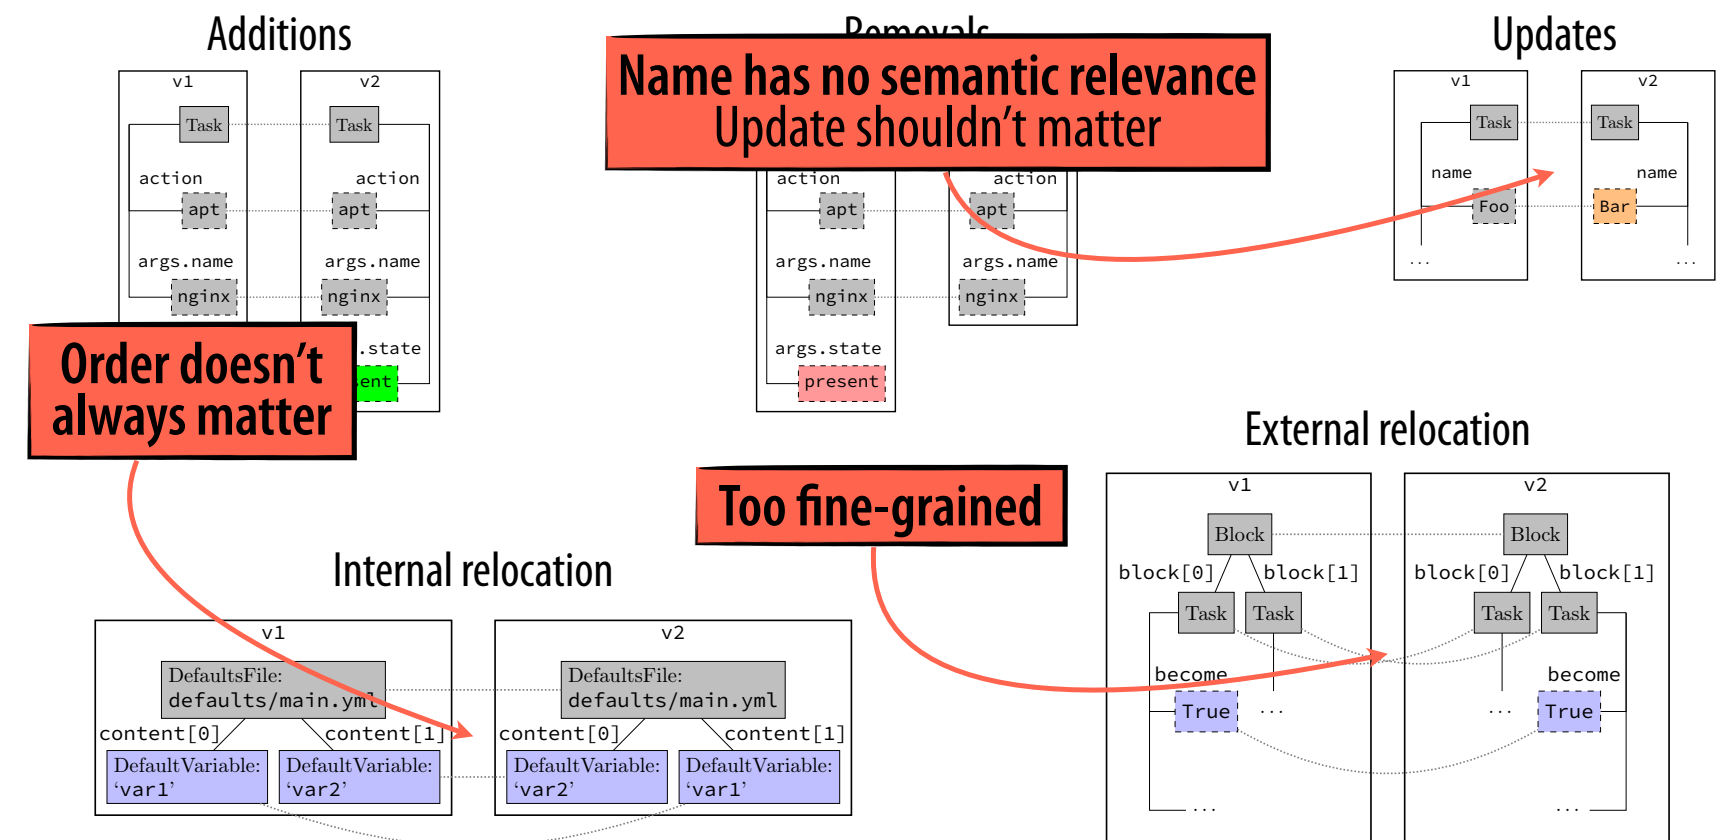
Conclusion

Structural Representation of Ansible Roles



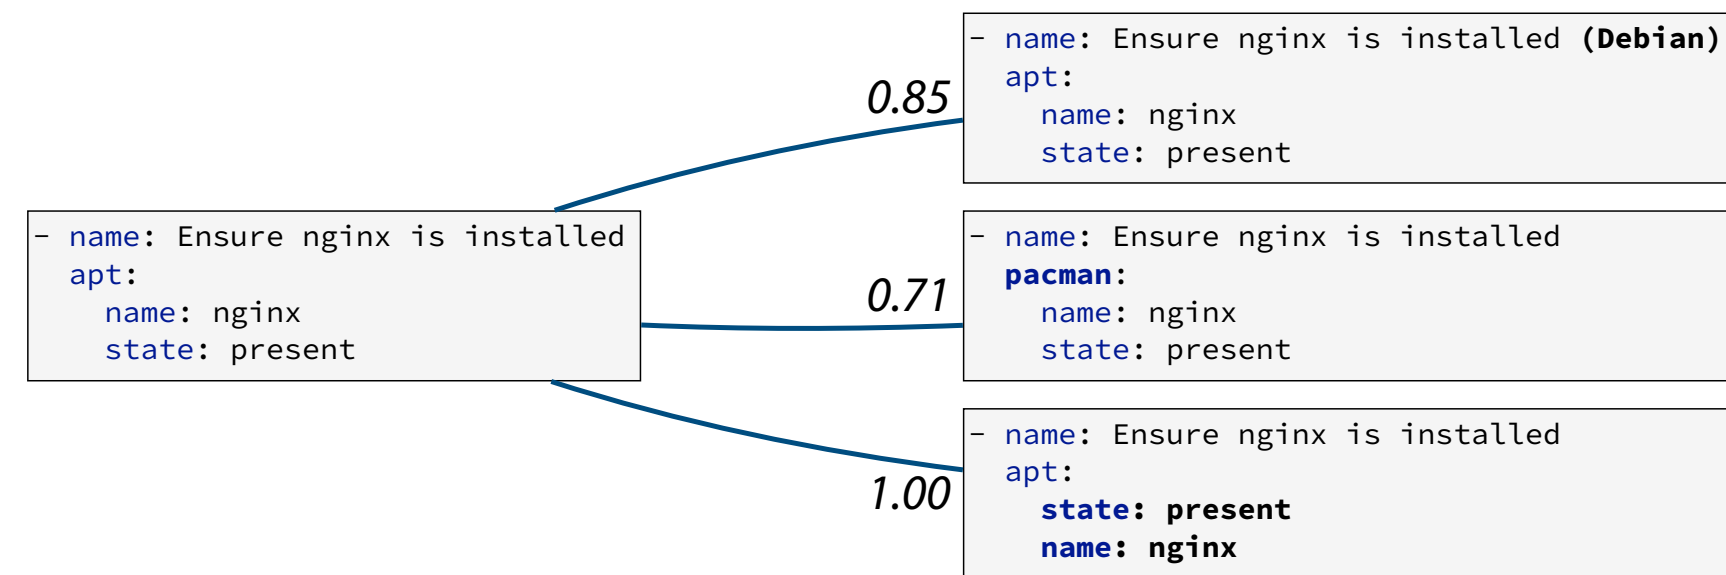
5

Change Distilling of Ansible Roles



8

Task Similarity



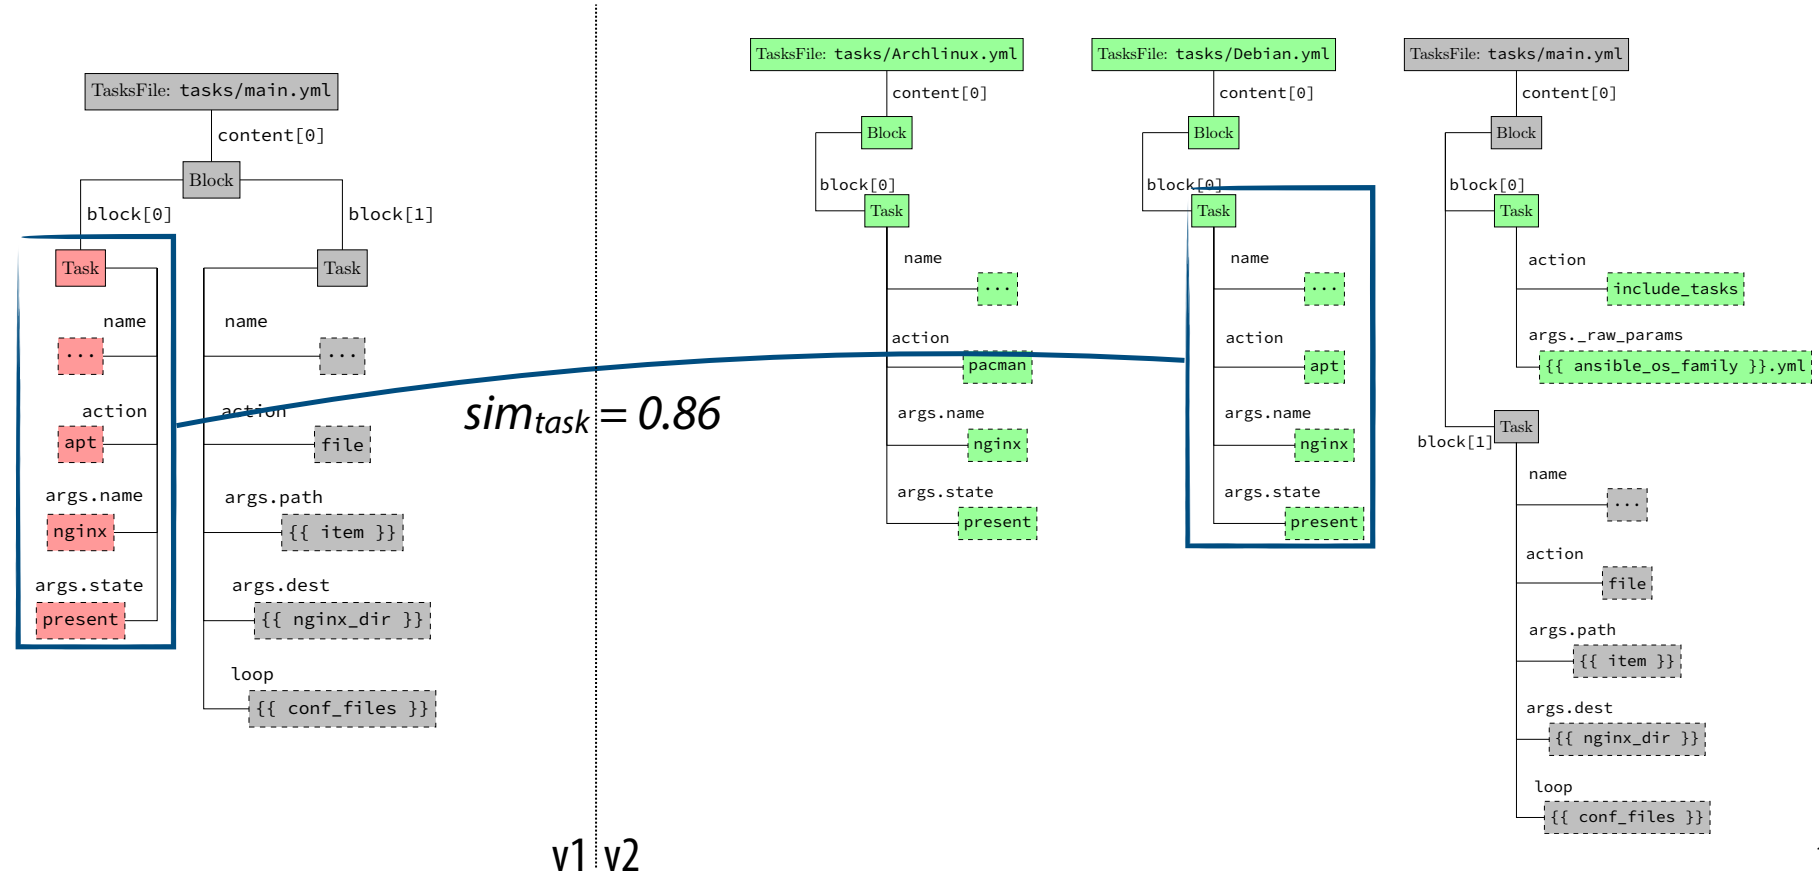
$$sim_{task}(T_1, T_2) = \frac{\sum_{kw \in T_1 \cap T_2} w(kw)}{\sum_{kw \in T_1 \cup T_2} w(kw)}$$

$$w(kw) = \begin{cases} 1, & \text{if } kw \in \{\text{action, args, when, loop, loop_control}\} \\ 0.5, & \text{otherwise} \end{cases}$$

"Weighted Jaccard"

9

Task Relocation Matching



v1 v2

11