

Configuration Management

Markus Raab

Institute of Information Systems Engineering, TU Wien

08.05.2018



Lecture is every week Wednesday 09:00 - 11:00.

06.03.2019: topic, teams

13.03.2019: TISS registration, initial PR

20.03.2019: other registrations, guest lecture

27.03.2019: PR for first issue done, second started

03.04.2019: first issue done, PR for second

10.04.2019: mid-term submission of exercises

08.05.2019: different location: **Complang Library**

15.05.2019:

22.05.2019: all 5 issues done

29.05.2019:

05.06.2019: final submission of exercises

12.06.2019:

19.06.2019: last corrections of exercises

26.06.2019: exam

Tasks for today

(until 08.05.2019 23:59)

Task

Incorporate feedback for teamwork and homework.

Task

Third PR done, PR for fourth issue created and write some text in your last issue (if 5 issues are not yet assigned to you).

Tasks for next week

(until 15.05.2019 23:59)

Task

Fourth PR done, PR for fifth issue created.

Task

Continue teamwork and homework.

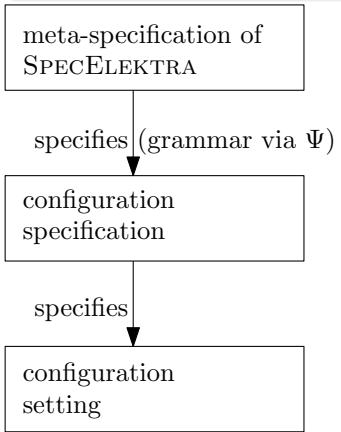
Popular Topics

- 14 tools
- 9 testability
- 9 code-generation
- 7 context-awareness
- 6 specification
- 6 misconfiguration
- 6 complexity reduction
- 5 validation
- 5 points in time
- 5 error messages
- 5 auto-detection
- 4 user interface
- 4 introspection
- 4 design
- 4 cascading
- 4 architecture of access
- 3 configuration sources
- 3 config-less systems
- 2 secure conf
- 2 architectural decisions
- 1 push vs. pull
- 1 infrastructure as code
- 1 full vs. partial
- 1 convention over conf
- 1 CI/CD
- 0 documentation

Metalevels (Recapitulation)

Question

Describe the three Metalevels in Elektra.



SpecElektra (Recapitulation)

SpecElektra is a modular ***configuration specification language*** for configuration settings. In SpecElektra we use properties to specify configuration settings and configuration access. SpecElektra enables us to specify different parts of Elektra.

Recapitulation (Requirements of SpecElektra)

- formal and informal
- should strive for completeness
- should be extensible
- should be external to application
- open for introspection (for tooling)
- should talk to users
- should allow generation of artefacts

Modularity (Recapitulation)

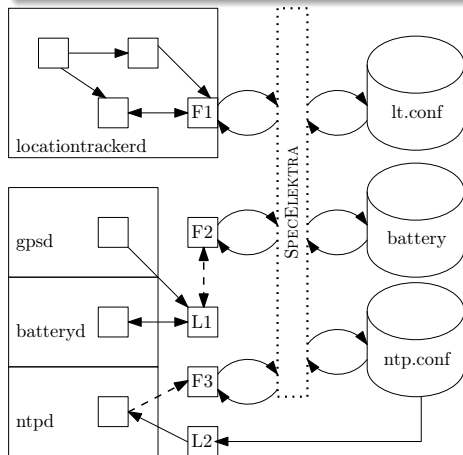
Vertical modularity describes how strongly separated the configuration accesses of different applications is.

Horizontal modularity describes how strongly separated modules implementing configuration access for a single application is.

Vertical Modularity (Recapitulation)

Question

Explain the content of the figure.



Plugins (Recapitulation)

Plugins are filters, sinks, and sources processing a key set. We aim at SpecElektra to be as modular as possible and make extensive use of plugins:

- 1 SpecElektra does not have any built-in feature, all features are (or can be) implemented as plugins.
- 2 Elektra works completely without SpecElektra's specifications.
- 3 Configuration specifications are present within the execution environment. Thus any tool and plugin can introspect and use the specifications.

Introspection (Recapitulation)

- unified get/set access to (meta*)-key/values
- access via applications, CLI, GUI, web-UI, ...
- access via any programming language (similar to file systems)
- GUI, web-UI can semantically interpret metadata

Goals for today

learning outcome:

- evaluate a configuration system and decide about use of
 - code generation
 - system-wide introspection

Code Generation

- 1 Code Generation
 - Why?
 - How?
- 2 Introspection vs. Generation

Task

How to ensure that configuration access points match with present configuration settings?

Rationale (Partly Recapitulation)

Configuration Specification:

- without specification you and others do not even know which settings are available
- needed for any further techniques we will discuss:
 - code generation guarantees that configuration access points match with specification
 - validation guarantees that configuration settings match with specification
- essential for *no-futz computing* Holland et al. [1]
- the foundation for any advanced tooling like configuration management tools
- needed as communication of producers and consumers of configuration

Why?

Current Challenges

Configuration access code usually has:

- code duplications and unsafe APIs
- hard-coded default values
- unexpected transformations (e.g., truncating of values)
- inconsistencies (e.g., case sensitivity)
- no introspection facilities (which keys and values are allowed?)

Example (Silent Overruling [4])

```
1 if (!strcasecmp(token, "on")) {
2     *var = 1;
3 } else {
4     *var = 0;
5 } /* src/cache_cf.cc from Squid */
```

Real-world example

PostgreSQL¹ has following duplications for its configuration settings:

- a global variable and an option record (struct)
- an entry in an example (postgresql.conf.sample)
- documentation in sgml
- in the source code of utils (in-source dump utils, and dozens of external configuration management tools)

Note: PostgreSQL has a clean implementation, and above list only shows limitations of systems without code generation.

¹http://doxygen.postgresql.org/guc_8c_source.html

Task

Brainstorming: Which artefacts can we produce with (code) generation?

Artefacts:

- examples (e.g., defaults)
- documentation
- auto-completion/syntax highlighting/IDE support
- tooling (GUI, Web UI)
- validation code
- configuration management tool code
- configuration access APIs

Why?

Goal

Goal

Configuration settings should adhere the specification from source to destination.

Requirement

The specification must enable code generation and inconsistencies must be ruled out during compilation.

Code Generation

The code generator GenElektra reads SpecElektra specifications and emits high-level APIs to be used in applications. GenElektra facilitates the key names to generate unique API names.

But how?

How?


KeySet (Recapitulation)

The common data structure between plugins:

KeySet

Key

metakey or property



How?

KeySet Generation

Question

Idea: What if the configuration file format grammar describes source code?

$$\langle \textit{KeySet} \rangle ::= \textit{'ksNew' } \sqcup (\{ \langle \textit{Key} \rangle \textit{' , } \leftarrow \} \{ \textit{' } \sqcup \} \textit{'KS_END'});$$

$$\langle \textit{Key} \rangle ::= \textit{'keyNew' } \sqcup (\textit{' ' } \langle \textit{key name} \rangle \textit{' ' } \textit{' , } \leftarrow [\langle \textit{Value} \rangle]$$

$$\langle \textit{properties} \rangle \textit{'KEY_END'}$$

$$\langle \textit{Value} \rangle ::= \{ \textit{' } \sqcup \} \textit{'KEY_VALUE, } \sqcup \textit{' ' } \langle \textit{configuration value} \rangle \textit{' ' } ,$$

$$\leftarrow$$

$$\langle \textit{properties} \rangle ::= \{ \{ \textit{' } \sqcup \} \langle \textit{property} \rangle \textit{' , } \leftarrow \}$$

$$\langle \textit{property} \rangle ::= \textit{'KEY_META, } \sqcup \textit{' ' } \langle \textit{property name} \rangle \textit{' ' } ,$$

$$\sqcup \textit{' ' } \langle \textit{property value} \rangle \textit{' '}$$

How?

Task

Break.

How?

Example

Example

Given the key spec `/slapd/threads/listener`, with the configuration value 4 and the property default \mapsto 1, GenElektra emits:

```
1 ksNew (keyNew ("spec:/slapd/threads/listener",
2             KEY_VALUE, "4",
3             KEY_META, "default", "1",
4             KEY_END),
5         KS_END);
```

Finding

We have source code representing the settings. And if we instantiate it, we have a data structure representing the settings. Plugins emitting such “configuration files” are code generators.

Implementation Strategies

- Using print (only for very small generators)
- Using generative grammars

```

1 query = '{' >> *(pair) > '}';
2 pair = '{' >> key_name > '=' >> key_value >>
3       *('{ ' >> metakey_name > '=' >> metakey_
4       > '}');
```

- Using template languages (RubyERB, Cheetah, Mustache)

```

1 @for n in hierarchy.name.split('/')[1:-1]
2 namespace $support.nsnpretty($n)
3 {
4 class ${hierarchy.prettyclassname(support)}
5 {
6 typedef $support.typeof($hierarchy.info) type;
7 @if $support.typeof($hierarchy.info) != "kdb::none_t"
8 static type get(kdb::KeySet &ks, kdb::Key const& spec)
9 {
10     type value $support.valof($hierarchy.info)
11     Key found(ckdb::ksLookup(ks.getKeySet(), *spec,
12                          ckdb::elektraLookupOptions::KDB_0_SPEC));
13     return found.get<$support.typeof($hierarchy.info)>();
14 }
```

Possible Properties

For example, SpecElektra has following properties:

- type** represents the type to be used in the emitted source code.
- opt** is used for short command-line options to be copied to the namespace proc.
- opt/long** is used for long command-line options, which differ from short command-line options by supporting strings and not only characters.
- readonly** yields compilation errors when developers assign a value to a contextual value within the program.
- default** enables us to start the application even if the backend does not work.

With the specification:

```
1 [foo/bar]
2   default := Hello
3   type := string
4   opt := b
5   readonly := 1
```

GenElektra gives the user read-only access to the object `env.foo.bar`:

```
1   std::cout << env.foo.bar;
2   env.foo.bar = "Other world"; // comp. error
```

Line 1 prints the configuration value of `/foo/bar` or `"Hello"` (without quotes) by default. When invoking the application with `application -b "This world"`, the application would print `"This world"` (without quotes). Line 2 leads to a compilation error because of the property `readonly`.

Which Configuration Access API?

First approach, one class (or function) per configuration setting:

```
1 class SlapdThreadsListener : public Value<long ,
2     WritePolicyIs<ReadOnlyPolicy>> {
3     ... keyNew ("/slapd/threads/listener" ,
4         KEY_META , "type" , "long" ,
5         KEY_META , "readonly" , "1" ,
6         KEY_END) ...
7 };
```

Which Configuration Access API?

Bad idea, manual instantiation and long names necessary:

```
1 KeySet config;
2 Context c;
3 long foo ()
4 {
5     SlapdThreadsListener slapdThreadsListener (con
6     slapdThreadsListener++;
7     return slapdThreadsListener;
8 }
```

How?

Which Configuration Access API?

Use hierarchy with namespaces or nested classes:

```
1 namespace slapd
2 {
3 namespace threads
4 {
5 class Listener : public Value<long> {};
6 } // <continues on the next page>
7 class Threads : public Value<none_t>
8 {threads::Listener listener;};
9 } // end namespace slapd
10 class Slapd : public Value<none_t>
11 {slapd::Threads threads;};
12 class Environment {Slapd slapd;};
```


How?

Which Configuration Access API?

Much easier to use:

```
1 long foo(slapd::Threads const & threads)
2 {
3     threads.listener++;
4     Context & c = threads.context ();
5     return threads.listener;
6 }
7
8 int main()
9 {
10    KeySet config;
11    Context c;
12    Environment env (config, c);
13    long x = foo (env.slapd.threads);
14 }
```

Which Configuration Access API?

In C, we use identifiers to be passed to the highlevel API¹:

```
1 elektraGetString (elektra, ELEKTRA_TAG_MY);
```

Where `ELEKTRA_TAG_MY` is a struct for that type.

We can also omit the type:

```
1 elektraGetLong (elektra, ELEKTRA_TAG_THREADS);
```

```
2 elektraGet (elektra, ELEKTRA_TAG_THREADS);
```

¹<https://www.libelektra.org/tutorials/high-level-api>

Guarantees by code generation:

- Every configuration setting is specified.
- (Data) type of source code and configuration settings match.
- Configuration access with defaults is always successful.
Reason: We use defaults if everything else fails.

Missing Guarantee: Is every specified setting actually used?

Introspection vs. Generation

- 1 Code Generation
 - Why?
 - How?
- 2 Introspection vs. Generation

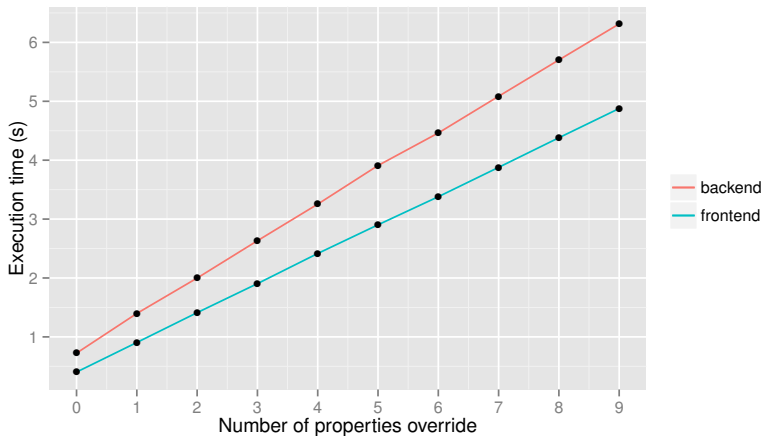
Question

Introspection vs. Code Generation?

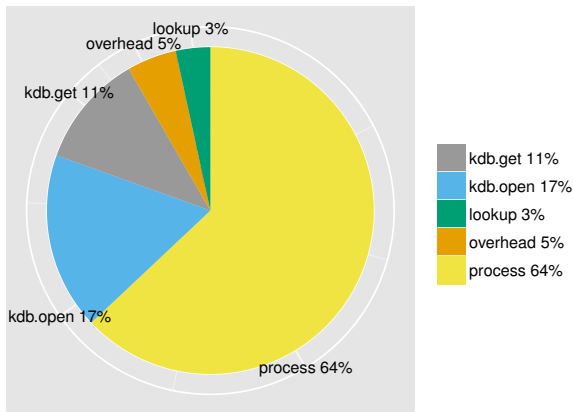
Limitations of introspection:

- no static checks
- no whole-program optimizations (API barriers)

Overhead without code generation (=backend) is 1.8x higher [2]:



But it might not matter because configuration access might not be a bottleneck [2], for example, a word counting application:



But: Configuration access points within loops might be a bottleneck.

Advantages of introspection:

- specification can be updated live on the system without recompilation
- tooling has generic access to all specifications
- new features the key database (e.g., better validation) are immediately available consistently

Implication

We generally prefer introspection, except for a very thin configuration access API.

Requirement

Configuration settings and specifications must be introspectable.

Preview

- Testing
- Early Detection of Misconfiguration

- [1] David A. Holland, William Josephson, Kostas Magoutis, Margo I. Seltzer, Christopher A. Stein, and Ada Lim. Research issues in no-futz computing. In *Hot Topics in Operating Systems, 2001. Proceedings of the Eighth Workshop on*, pages 106–110. IEEE, May 2001. doi: 10.1109/HOTOS.2001.990069.
- [2] Markus Raab. Sharing software configuration via specified links and transformation rules. In *Technical Report from KPS 2015*, volume 18. Vienna University of Technology, Complang Group, 2015.
- [3] Markus Raab and Gergö Barany. Introducing context awareness in unmodified, context-unaware software. In *Proceedings of the 12th International Conference on Evaluation of Novel Approaches to Software Engineering - Volume 1: ENASE,,* pages 218–225. INSTICC, ScitePress, 2017. ISBN 978-989-758-250-9. doi: 10.5220/0006326602180225.

- [4] Tianyin Xu, Jiaqi Zhang, Peng Huang, Jing Zheng, Tianwei Sheng, Ding Yuan, Yuanyuan Zhou, and Shankar Pasupathy. Do not blame users for misconfigurations. In *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles*, pages 244–259. ACM, 2013.