






The Pain and Pleasure of Building a Database Engine

LARS RÖNNBÄCK
The Knowledge Gap 2022





Java, C++, C, Scala, Go, Haskell,
Swift, .NET, Rust, Node.js, WASM,
Python, Ruby, SQL, Perl, Pascal,
Basic, Kotlin, Julia, Lisp, Elm...

I need the
BEST
possible
language.



But how do you
decide? How do you
find good and reliable
snippets? How do you
know if it's "alive"?

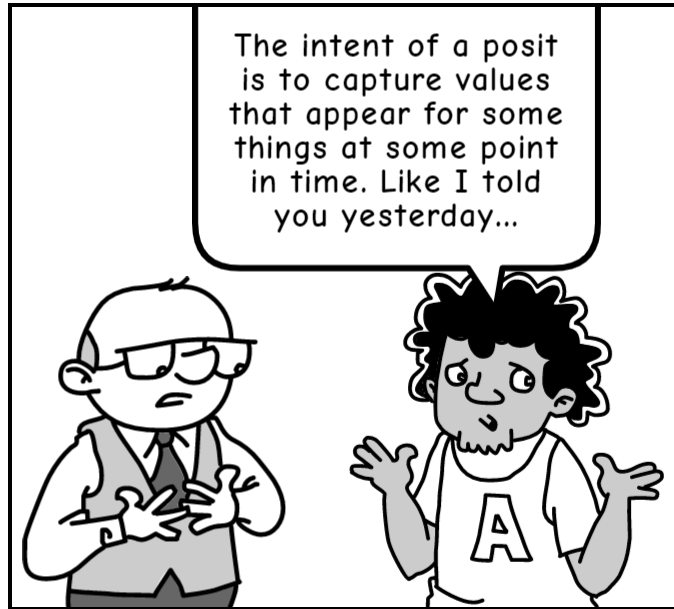
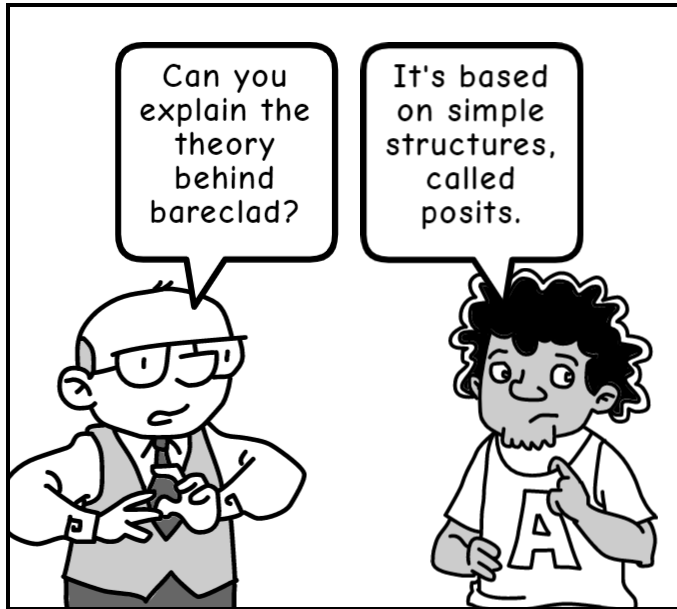


I know, I'll use the
CHOAS metric:

Clueless
Has
Opinion
About
Subject



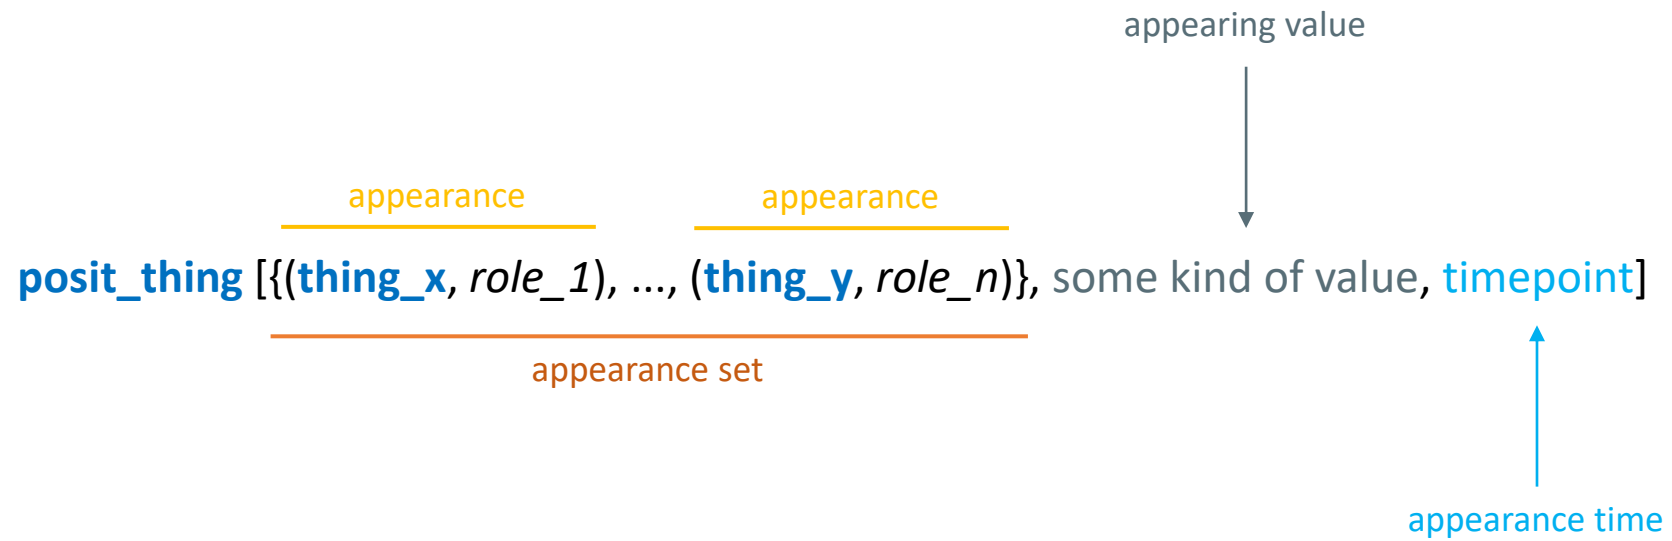
...and Rust is the winner!



posit_thing [{(thing_x, role_1), ..., (thing_y, role_n)}, some kind of value, timepoint]

posit – an atomic data piece

posit_thing [{(thing_x, role_1), ..., (thing_y, role_n)}, some kind of value, timepoint]




```
p1 [{(idh, husband), (idw, wife)}, "married", '2004-06-19']
```

The same **thing** can appear in many posits

The same **role** can appear in many posits

The same **appearance** can appear in many posits

p1 [{(idh, husband), (idw, wife)}, "married", '2004-06-19']

The same **appearance set** can appear in many posits

The same **appearance time** can appear in many posits

The same appearing value can appear in many posits

MANY WAYS TO SEARCH!

All posits in which this **thing** appears

All posits in which this *role* appears

All posits in which this **appearance** appears


p1 [{(idh, husband), (idw, wife)}, "married", '2004-06-19']

All posits in which this **appearance set** appears

All posits in which this **appearance time** appears

All posits in which this appearing value appears

Traqula



add role *wife, husband, name, age, address, epithet*;

add posit

```
+p1 [{(+idw, wife), (+idh, husband)}, "married", '2004-06-19'],
+p1 [{(idh, name)}, "Lars Samuelsson", '1972-08-20'], /* A name I used to go by */
+p1 [{(idh, address)}, {
    "street": "The Road of Karlberg",
    "postal code": 11335,
    "care_of": {
        "name": "Lars Roenbaeck",
        "apartment": 1501
    }
}, '2007-09-01'],
+p1 [{(idh, age)}, 3.14159265358979323846264338327950288419716939937510582097494459, '2003-02-20'],
+p2 [{(idh, epithet)}, "/* liker of comments */", ], and ""spaces"", @BOT];
```

add role *wife, husband, name, age, address, epithet;*

add posit

```
+p1 [{(+idw, wife), (+idh, husband)}, "married", '2004-06-19'],
+p1 [{(idh, name)}, "Lars Samuelsson", '1972-08-20'], /* A name I used to go by */
+p1 [{(idh, address)}, {
  "street": "The Road of Karlberg",
  "postal code": 11335,
  "care_of": {
    "name": "Lars Roenbaeck",
    "apartment": 1501
  }
}, '2007-09-01'],
+p1 [{(idh, age)}, 3.14159265358979323846264338327950288419716939937510582097494459, '2003-02-20'],
+p2 [{(idh, epithet)}, "/* liker of comments */", ], and ""spaces"", @BOT];
```

roles

add role *wife, husband, name, age, address, epithet*;

add posit

```
+p1 [{(+idw, wife), (+idh, husband)}, "married", '2004-06-19'],  
+p1 [{(idh, name)}, "Lars Samuelsson", '1972-08-20'], /* A name I used to go by */  
+p1 [{(idh, address)}, {  
  "street": "The Road of Karlberg",  
  "postal code": 11335,  
  "care_of": {  
    "name": "Lars Roenbaeck",  
    "apartment": 1501  
  }  
}, '2007-09-01'],  
+p1 [{(idh, age)}, 3.14159265358979323846264338327950288419716939937510582097494459, '2003-02-20'],  
+p2 [{(idh, epithet)}, "/* liker of comments */", ], and ""spaces"", @BOT];
```


A diagram illustrating the use of the `things` variable. A green box labeled `things` has two green arrows pointing to the `idh` variables in the third and fourth lines of the code block. The first arrow points to the `idh` in `[(idh, address)]`, and the second arrow points to the `idh` in `[(idh, name)]`.

add role *wife, husband, name, age, address, epithet*;

add posit

```
+p1 [{(+idw, wife), (+idh, husband)}, "married", '2004-06-19'],
+p1 [{(idh, name)}, "Lars Samuelsson", '1972-08-20'], /* A name I used to go by */
+p1 [{(idh, address)}, {
  "street": "The Road of Karlberg",
  "postal code": 11335,
  "care_of": {
    "name": "Lars Roenbaeck",
    "apartment": 1501
  }
}, '2007-09-01'],
+p1 [{(idh, age)}, 3.14159265358979323846264338327950288419716939937510582097494459, '2003-02-20'],
+p2 [{(idh, epithet)}, "/* liker of comments */", ], and ""spaces"", @BOT];
```

timepoints

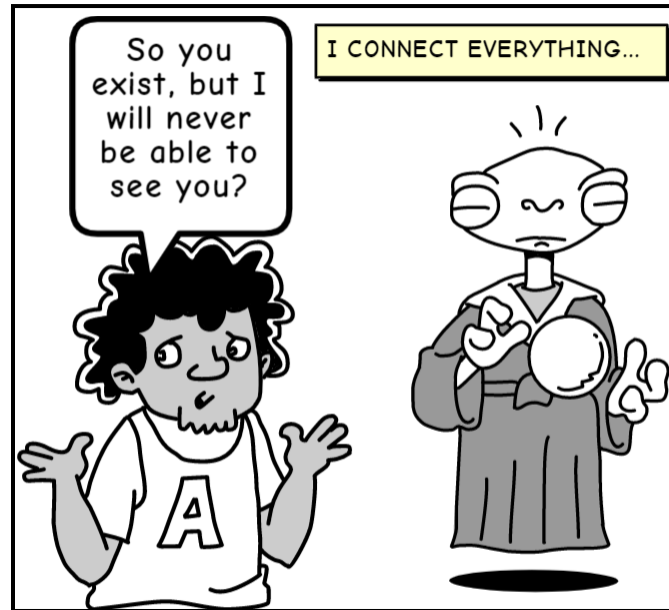



add role *wife, husband, name, age, address, epithet*;

add posit

```
+p1 [{(+idw, wife), (+idh, husband)}, "married", '2004-06-19'],  
+p1 [{(idh, name)}, "Lars Samuelsson", '1972-08-20'], /* A name I used to go by */  
+p1 [{(idh, address)}, {  
    "street": "The Road of Karlberg",  
    "postal code": 11335,  
    "care_of": {  
        "name": "Lars Roenbaeck",  
        "apartment": 1501  
    }  
}, '2007-09-01'],  
+p1 [{(idh, age)}, 3.14159265358979323846264338327950288419716939937510582097494459, '2003-02-20'],  
+p2 [{(idh, epithet)}, "/* liker of comments */", ], and ""spaces"", @BOT];
```

A diagram illustrating the concept of 'values' in the context of the provided code. A green box labeled 'values' has four arrows pointing to specific values in the code: 'married', '1972-08-20', '2007-09-01', and '2003-02-20'. These values are highlighted in blue in the original image.



add role *wife, husband, name, age, address, epithet*;

add posit insert

+**p1** [{(**+idw**, wife), (**+idh**, husband)}, "married", '2004-06-19'],
+**p1** [{(**idh**, name)}, "Lars Samuelsson", '1972-08-20'], /* A name I used to go by */
+**p1** [{(**idh**, address)}, {
 "street": "The Road of Karlberg",
 "postal code": 11335,
 "care_of": {
 "name": "Lars Roenbaeck",
 "apartment": 1501
 }
}], '2007-09-01'],
+**p1** [{(**idh**, age)}, 3.14159265358979323846264338327950288419716939937510582097494459, '2003-02-20'],
+**p2** [{(**idh**, epithet)}, "/* liker of comments */",], and ""spaces"", @BOT];

recall

```
graph TD
    insert[insert] --> add_posit[add posit]
    insert --> idw((+idw))
    recall[recall] --> idh2((idh))
```

add role *birth date*;

add posit

```
[{(idw, name)}, "Anneli", '1972-02-13'],  
[{(idw, age)}, 30, '2002-02-13'],  
[{(idw, birth date)}, '1972-02-13', '1972-02-13'];
```

add posit +p1 [{(idh, name)}, "Lars Roenbaeck", '2004-07-01'];

add posit

```
[{(p1, posit), (idh, ascertains)}, 100%, @NOW],  
[{(p2, posit), (idh, ascertains)}, -100%, @NOW];
```

add posit

```
[{(+c, name)}, "Person", @NOW],  
[{(idw, thing), (c, class)}, "belongs", '1972-02-13'];
```

add role *birth date*;

add posit

```
[{(idw, name)}, "Anneli", '1972-02-13'],  
[{(idw, age)}, 30, '2002-02-13'],  
[{(idw, birth date)}, '1972-02-13', '1972-02-13'];
```

add posit +**p1** [{(idh, name)}, "Lars Roenbaeck", '2004-07-01']; ← **p1** holds five posit things here

add posit

```
[{(p1, posit), (idh, ascertains)}, 100%, @NOW], ← so this expands to five posits  
[{(p2, posit), (idh, ascertains)}, -100%, @NOW];
```

add posit

```
[{(+c, name)}, "Person", @NOW],  
[{(idw, thing), (c, class)}, "belongs", '1972-02-13'];
```

add role *birth date*;

add posit

```
[{(idw, name)}, "Anneli", '1972-02-13'],  
[{(idw, age)}, 30, '2002-02-13'],  
[{(idw, birth date)}, '1972-02-13', '1972-02-13'];
```

add posit +**p1** [{(idh, name)}, "Lars Roenbaeck", '2004-07-01'];

add posit

```
[{(p1, posit), (idh, ascertains)}, 100%, @NOW],  
[{(p2, posit), (idh, ascertains)}, -100%, @NOW];
```


reserved roles | assertions (metadata)

add posit

```
[{(+c, name)}, "Person", @NOW],  
[{(idw, thing), (c, class)}, "belongs", '1972-02-13'];
```

reserved roles | classification (peridata, i.e. peripheral data)

SEARCH



First you specify what patterns you want to match with your query.



WHERE

Then you limit the matches using variables that were populated in the search.



RETURN

Finally, you list what of the matched content you want in your result.



search

```
[[(idw, wife), (*, husband)], "married", t] as of @NOW,  
[[(idw, name)], w, *] as of t
```

where

```
t >= '2001-01-01'
```

aggregate

```
for w { number_of_wives_with_name = count(idw) }
```

return

```
w,  
number_of_wives_with_name;
```

And this is the parser
I am currently
working on and expect
to finish before the
summer. After that
there will be an alpha
release, but still lots
to be done!



Peril	bareclad	relational database	What do you mean by that?
Boundaries	🤝	🤝	Where a thing begins and ends.
Vocabulary	🤝	🤝	What words we use to talk about things.
Classification	✓	🤝	How we categorize things that have similarities.
Intrinsics	✓	🤝	How values are represented (data types).
Uncertainty	✓	🤝	How certain we are of the statements we have.
Imprecision	⌚	✓	Data types that can store imprecise values.
Identification	⌚	✗	Search criteria for recollecting a thing.
Evolution	✓	✗	Everything can change over time.
Fallibility	✓	✗	Corrections may need to be made.
Negation	✓	✗	Talking about the opposite of a statement.
Relativity	✓	✗	Conflicting opinions are not unheard of.
Evaluation	⚠️	⚠️	A value is subjectively determined on the way in.
Interpolation	⚠️	⚠️	A value is subjectively experienced on its way out.
Insufficient context	⚠️	⚠️	A situation can never be captured in its entirety.

- 🤝 must be universally agreed upon
- ✓ managed by the database engine
- ✗ not managed by the database engine
- ⚠️ cannot be managed by a database engine
- ⌚ upcoming feature

