

# Introduction to Relational Database Systems

## Datenbanksysteme 1 (INF 3131)

**Torsten Grust**  
**Universität Tübingen**  
Winter 2013/14

# Typed / Untyped Data

## Valid Data?

1. In the **text data model**, any `'\n'`-separated file of (UTF8-encoded) text represents valid data.  
(All we can hope for is that the different lines contain the same patterns of text.)
2. For **JSON**, any value derivable from the non-terminal symbol `<value>` is valid data.
3. For the **tabular data model**, any CSV file in which each row (including the header row) contains the same number of `'\t'`-separated fields is valid data.

## Valid, But Loose

The above are rather loose syntactic restrictions. It is still (too) easy to craft valid data that makes queries trip. Query execution may halt or (even worse) return non-sensical results.

# Typed / Untyped Data

## Example: Mangle the LEGO Set 5610 Text Data

- Slight edit in minifigure weight data (3.27g → 0.3.27g):

```
...  
Minifig#  Weight  Name  
1x cty052    0.3.27g  Construction Worker ...  ⚠
```

- Line ⚠ still matches the `sed` regular expression `^([0-9]+)x.[ /]([0-9.]+)g.*$`
- Rules of `awk`'s “best effort” string-to-number conversion apply.  
Overall result (weight of set) misleading.
- This is a so-called **silent error**. Particularly hard to detect in a sea of data.

# Typed / Untyped Data

- JSON arrays and dictionaries may have **heterogeneous contents** (any `<value>` is like any other):

```
let $xs := [1, 2, "three", 4, 5]    ⚠  
return  
  for $x in members($xs)  
  return $x + 1
```

- JSONiq query will fail at **query runtime** (i.e., rather late; NB: query syntax is okay):

```
Error  
[err:XPTY0004] arithmetic operation not defined between types "xs:string"  
and "xs:integer" at line 4, column 12
```

# Typed / Untyped Data

- In JSONiq, failing index or key lookups **silently** evaluate to `()` (empty sequence):

```
{ "one": 1 }. "one"      → 1
{ "one": 1 }. "two"     → ()
[1, 2, 3][[4]]         → ()
```

- Expressions involving `()` **propagate** `()`, which makes debugging of JSONiq queries hard:

```
() + 1                → ()
{ "one": 1 }. ()      → ()
[1, 2, 3][()]         → ()
for $x in members([1,2,3]) where () return $x → ()
```

# Typed / Untyped Data

- In PyQL and tabular CSV data, we rely on certain fields to contain data of a certain **type** (e.g., number or Boolean).
- If these assumptions are wrong, explicit conversions like `int()` or `float()` may fail at **query runtime** (after many seconds/minutes/hours, if we are unlucky).
- “Safe conversions” can help here but may introduce a noticeable runtime overhead:

```
# convert string s to float (if that fails, return default float x instead)
def safe_float(s, x):
    try:
        f = float(s)
    except ValueError:
        f = x

    return f
```

# Typed Data in the Relational Data Model

## Untyped Data Models

The text, JSON, and the tabular data models do *not* enforce (container or atomic) values to be of specific types. These data models are thus referred to as being **untyped**.

- The **relational data model** may be understood as a **typed** variant of the tabular data model:
  1. There is only one container type: **table** (or: multisets) **of rows**,
  2. all rows are of the **same row type** which is declared when the table is created,
  3. A row type **consists** of a sequence of **atomic types**.

# Typed Data in the Relational Data Model

- **SQL:** Creating a table  $\langle t \rangle$  and declaring its **row type**:

```
CREATE TABLE  $\langle t \rangle$  (                                -- table name and type name
     $\langle col_1 \rangle$   $\langle ty_1 \rangle$ ,
    ⋮
     $\langle col_r \rangle$   $\langle ty_r \rangle$ 
);
```

- Creates table  $\langle t \rangle$  of  $r$  columns,  $\langle col_i \rangle$  column name, *all* values in that column of atomic type  $\langle ty_i \rangle$ .
- Also implicitly declares row type  $\langle t \rangle = (\langle col_1 \rangle \langle ty_1 \rangle, \dots, \langle col_r \rangle \langle ty_r \rangle)$ .



# Typed Data in the Relational Data Model

- **Import** (or: load) correctly typed data from CSV file into a relational table (here: Relational Database Management System PostgreSQL):

```
COPY <t>(<col1>, ..., <colr>) FROM <csv>;
```

- Table <t> must have been created prior to the `COPY` command
- CSV file <csv> must *not* contain a header row: all rows are interpreted as data rows (⚠ specify an absolute path name for <csv> or use `stdin` to read from standard input)
- Order and number of fields in the CSV file must match the columns <col<sub>1</sub>>, ..., <col<sub>r</sub>>
- The **field values in the CSV file must have the declared types** <ty<sub>1</sub>>, ..., <ty<sub>r</sub>>

# Typed Data in the Relational Data Model

- Queries in the relational data model may rely on all rows having a known row type  
→ **No expensive runtime checks and conversions** with safety measures (ala `safe_float()`).
- Since types are known when the query is analyzed/compiled, the system may **specialize the executable query code** to (only) work for those types.  
(Even if we only save few milliseconds per row, the savings add up if we process many rows — as is typical.)
- Type-based errors in queries are detected at **query compile time** and thus early.  
Once a query has been compiled successfully, **no type errors will occur at runtime.**

# Declarativity

- Once data instances (text files, JSON values, CSV files, relational tables) are of significant size and queries are of moderate complexity and beyond, **query performance** becomes a concern.
- Performance gains may be achievable in different ways:
  1. Carefully **exploit properties of the data** (mini-world rules) to simplify your queries.
  2. Find entirely **different querying strategies** (beyond naive nested iteration, say) to process the data.
- ⚠ Both options involve query modifications or whole query rewrites — **query equivalence** may *not* be sacrificed.
- Let us see whether such performance gains (here: **reduction of elapsed query time**) are achievable for the *weight of LEGO Set 5610* query.

# Declarativity

- **Baseline:** Original PyQL program for the *weight of LEGO Set 5610* query:

```
from DB1 import Table

contains = Table('contains.csv')
bricks   = Table('bricks.csv')
minifigs = Table('minifigs.csv')

sum = 0

for c in contains:
    if c['set'] == '5610-1':
        for b in bricks:
            if c['piece'] == b['piece']:
                sum += int(c['quantity']) * float(b['weight'])
        for m in minifigs:
            if c['piece'] == m['piece']:
                sum += int(c['quantity']) * float(m['weight'])

print(sum)
```

# Declarativity

- Let  $|S|$  denote the *cardinality* of set  $S$   
(i.e.,  $|\langle t \rangle|$  denotes the number of rows in CSV table  $\langle t \rangle$ )
- Measure the work performed by the PyQL programs in terms of the **numbers of rows** processed.
  - The work per row (field access, arithmetics, comparison) is assumed to be essentially constant, i.e. in  $O(1)$ .
  - Let  $pieces(5610)$  denote the number of different piece types in LEGO Set 5610  
(a one-line PyQL comprehension can compute function  $pieces(s)$  for any given set  $s$ )

Work performed by baseline PyQL query:

$$|\text{contains}| + pieces(5610-1) \times (|\text{bricks}| + |\text{minifgs}|)$$

# Declarativity

## Optimization #1 (Exploit Properties of the Data)

- Observations:
  1. Field `piece` is **unique** in table `brick`: no two bricks share the same piece identifier. Once we have found a brick, we will not find another.  
(The same is true for table `minifigs`.)
  2. A given piece identifier  $p$  is **either** found in table `bricks` **or** in table `minifigs`. The `piece` field values of both tables are disjoint.  
(Question: What about the third option:  $p$  not present in any of the two tables?)
- **NB**: Both observations are consequences of **rules** in the LEGO sets mini-world.
- We can use this to optimize the baseline PyQL query. Recall: **query equivalence** must be preserved.

# Declarativity

Estimation of work performed by PyQL query after optimization #1:

$$pieces(5610-1) \times (b \times \frac{1}{2} \times |bricks| + (1 - b) \times (|bricks| + \frac{1}{2} \times |minifigs|))$$

where  $0 \leq b \leq 1$  is the fraction of bricks in a set ( $b = 0.95$  for LEGO Set 5610).

- In this estimate,
  - $\frac{1}{2} \times |bricks|$  describes the average effort to find a piece in table `bricks`,
  - $|bricks| + \frac{1}{2} \times |minifigs|$  describes the effort to miss a piece in `bricks` and then find it in `minifigs`.

# Declarativity

## Optimization #2 (Change of Query Strategy)

- Idea: Proceed in two phases:
  - Phase 1: Iterate over `contains` and build a **small temporary data structure** that maps pieces to their quantity in LEGO Set 5610. (Only include pieces in Set 5610!)
  - Phase 2: Iterate over `bricks` and `minifigs` **once** and check for each brick / minifig in the data structure. If present, update overall sum weight appropriately.
- The auxiliary data structure essentially implements a function `piece`  $\mapsto$  `quantity`. That function is partial (domain contains pieces of Set 5610 only).
- Can think of the mapping as a two-column table itself:

<code>piece</code>	<code>quantity</code>
$p_1$	$n_1$
$p_2$	$n_2$
$\vdots$	$\vdots$



# Declarativity

Estimation of work performed by PyQL query after optimization #2:

$$|\text{contains}| + |\text{bricks}| + |\text{minifigs}|$$

- It essential that lookups in the temporary data structure (a variant of an **index**) are low-cost:
  - Size of index in  $O(\text{pieces}(5610-1)) = O(1)$
  - → We can expect  $O(1)$  lookup

# Declarativity

- Database systems are designed to **shield** users from performance considerations of the type we have just studied:
  - The exploitation of data properties (e.g. uniqueness, disjointness) is automatic.
  - Decisions about the construction/removal of auxiliary data structures are either automatic or performed by database system administrators (DBAs).
  - Exploitation of such auxiliary data structures is automatic.

## Declarativity (“What, not how”)

Database systems offer **declarative query languages**. Users specify **what** results a query shall return, **not how** these results are to be computed efficiently.

- **SQL is declarative:** A SQL query expression *cannot* refer to auxiliary data structures that may be present for efficiency reasons. (There simply is no SQL query construct of this kind.)

# Data Independence

- In the LEGO sets mini-world, a *piece* is either a brick or a minifigure. All of our *Weight of LEGO Set 5610* queries were **explicitly** reflecting this fact. Cumbersome.
- Consider two options:
  1. In the PyQL queries, define a temporary `pieces` list that combines the common features of bricks and minifigures. Then query `pieces` instead.
  2. Construct a new `pieces` table (on disk) from the existing `bricks` and `minifigs` tables. In PyQL queries, refer to the new `pieces` table just like any other table.
- Both options sound reasonable but require different pre-processing steps. However, once pre-processing is complete, **queries using `pieces` need not know the difference.**

# Data Independence

## Pre-Processing for Option #1

- Perform pre-processing inside the PyQL program:-

```
⋮
bricks    = Table('bricks.csv')
minifigs  = Table('minifigs.csv')

# A piece is either a brick or a minifig: build the union (make sure to
# only retain the features common for both piece types)
pieces = [ m for m in minifigs ] \
        + \
        [ { 'piece':  b['piece'],
            'type':   b['type'],
            'name':   b['name'],
            'cat':    b['cat'],
            'weight': b['weight'],
            'img':    b['img'] } for b in bricks ]
⋮
```

# Data Independence

## Pre-Processing for Option #2

- Prepare a new table `pieces.csv` on disk. Use UNIX shell utilities to create the required CSV file:

```
> cut -f1-6 bricks.csv | tail +2 | cat minifigs.csv - > pieces.csv
```

- `cut -f1-6`: project on the first six (tab-separated) fields of each line
- `tail +2`: copy input to output but omit the first line (column name header)
- `cat`: concatenate inputs and copy to output (`-` denotes standard input)
- Inside the PyQL program:

```
⋮  
pieces = Table('pieces.csv')  
⋮
```

# Data Independence

- **Regardless of option chosen**, the rest of the PyQL query reads:

```
⋮
sum = 0

for c in contains:
    if c['set'] == '5610-1':
        for p in pieces:
            if c['piece'] == p['piece']:
                sum += int(c['quantity']) * float(p['weight'])

print(sum)
```

## Data Independence

The formulation of queries shall not depend on the details of physical data storage and layout. Customized views of the data are to be considered first-class data sources (just like tables).

# Persistence

- For a typical program, the output is a function of the (current user) input only:

- Today:

```
5 → factorial → 120
```

- Tomorrow:

```
5 → factorial → 120
```

- Program `factorial` does not consult data besides the user input and values/objects created by itself.

# Persistence

- For database-based programs (or: queries), the output is a function of the input **and the current state** of the mini-world (e.g., table contents):

- Today:

```
'5610-1'
contains.csv, bricks.csv, minifigs.csv ... } → weight of LEGO set → 22.46
```

- Tomorrow:

```
'5610-1'
contains.csv, bricks.csv, minifigs.csv ... } → weight of LEGO set → 25.30
```



# Persistence

## Persistence

Database systems maintain the state of data (tables) between invocations. Such *persistent* information outlives the process it was created by. In particular, the system persists the data across power outages and operating system reboots.

- Until now, we were using the **OS file system** to store persistent CSV files.
- This is where most database systems hold their persistent data, too.
  - ⚠ Plain file management via `open()`, `close()`, `read()`, `write()` is not sufficient. (Why?)
- Other storage options (e.g., raw disk devices) are common as well.