



CONTENTS

1	Data Tables and pandas	3
1.1	Data types	3
1.2	pandas	4
1.3	Reading a CSV with pandas	4
1.4	Looking at a Series	5
1.5	Rows and the index	6
1.6	Changing data	7
1.7	Derived columns	9
2	Data tables in SQL	11
2.1	Using SQL from Python	14
3	Diodes, Rectifiers, and Photovoltaics	17
3.1	Silicon	17
3.2	Diodes	19
3.2.1	Rectifiers	22
3.2.2	Photovoltaics	22
A	Answers to Exercises	25
	Index	27

Data Tables and pandas

Much of the data that you will encounter in your career will come to you as a table. Some of these tables are spreadsheets, some are in relational databases, and some will come to you as CSV files.

Typically, each column will represent an attribute (like height or acreage) and each row will represent an entity (like a person or a farm). You might get a table like this:

property_id	bedrooms	square_meters	estimated_value
7927	3	921.4	\$ 294,393
9329	2	829.1	\$ 207,420

In most cases, one of the columns is guaranteed to be unique. We call this the *primary key*. In this table, `property_id` is the primary key; every property has one, and no two properties have the same `property_id`.

1.1 Data types

Each column in a table has a type, and these usually correspond pretty nicely with types in Python.

	Type	Python type	Example
Here are some common datatypes:	Integer	<code>int</code>	910393
	Float	<code>float</code>	-23.19
	String	<code>string</code>	'Fred'
	Boolean	<code>bool</code>	False
	Date	<code>datetime.date</code>	2019-12-04
	Timestamps	<code>datetime.datetime</code>	2022-06-10T14:05:22Z

Sometimes it is OK to have values missing. For example, if you had a table of data about employees, maybe one of the columns would be `retirement` — a date that tells you when the person retired. People who had not yet retired would have no value in this column. We would say that they have *null* for `retirement`.

Sometimes there are constraints on what values can appear in the column. For example, if the column were `height`, it would make no sense to have a negative value.

Sometimes a column can only be one of a few values. For example, if you ran a bike rental shop, each bicycle's status would be "available", "rented", or "broken". Any other values

in that column would not be allowed. We often call these columns *categorical*.

1.2 pandas

The Python community works with tables of data *very* often, so it created the pandas library for reading, writing, and manipulating tables of data.

When working with tables, you sometimes need to go through them row-by-row. However, for large datasets, this is very slow. pandas makes it easy (and very fast) to say things like “Delete every row that doesn’t have a value for height” instead of requiring you to step through the whole table.

In pandas, there are two datatypes that you use a lot:

- a `Series` is a single column of data.
- a `DataFrame` is a table of data: it has a `Series` for each column.

In the digital resources, you will find `bikes.csv`. If you look at it in a text editor, it will start like this:

```
bike_id,brand,size,purchase_price,purchase_date,status
5636248,GT,57,277.99,1986-09-07,available
4156134,Giant,56,201.52,2005-01-09,rented
7971254,Cannondale,54,292.25,1978-02-28,available
3600023,Canyon,57,197.62,2007-02-15,broken
```

The first line is a header; it tells you the name of each column. Next, the each bike is represented by values which are separated by commas. (Thus the name: CSV stands for “Comma-Separated Values”.)

1.3 Reading a CSV with pandas

Let’s make a program that reads `bikes.csv` into a pandas dataframe. Create a file called `report.py` in the same folder as `bikes.csv`.

First, we will read in the csv file. pandas has one series that acts as the primary key; it calls this one the index. When reading in the file, we will tell it to use the `bike_id` as the index series.

If you ask a dataframe for its shape, it returns a tuple containing the number of rows and the number of columns. To confirm that we have actually read the data in, let’s print those

numbers. Add these lines to `report.py`:

```
import pandas as pd

# Read the CSV and create a dataframe
df = pd.read_csv('bikes.csv', index_col="bike_id")

# Show the shape of the dataframe
(row_count, col_count) = df.shape
print(f"*** Basics ***")
print(f"Bikes: {row_count:,}")
print(f"Columns: {col_count}")
```

Build it and run it. You should see something like this:

```
*** Basics ***
Bikes: 998
Columns: 5
```

Note that your table actually has six columns. The index series is not included in the shape.

1.4 Looking at a Series

Let's get the lowest, the highest, and the mean purchase price of the bikes. The purchase price is a series, and you can ask the dataframe for it. Add these lines to the end of your program:

```
# Purchase price stats
print("\n*** Purchase Price ***")
series = df["purchase_price"]
print(f"Lowest:{series.min()}")
print(f"Highest:{series.max()}")
print(f"Mean:{series.mean():.2f}")
```

Now when you run it, you will see a few additional lines:

```
*** Purchase Price ***
Lowest:107.37
Highest:377.7
Mean:249.01
```

What are all the brands of the bikes? Add a few more lines to your program that shows how many of each brand:

```
# Brand stats
print("\n*** Brands ***")
series = df["brand"]
series_counts = series.value_counts()
print(f"{series_counts}")
```

Now when you run it, your report will include the number of bikes for each brand from most to least common:

```
*** Brands ***
Canyon      192
BMC         173
Cannondale  170
Trek        166
GT          150
Giant       147
Name: brand, dtype: int64
```

`value_counts` returns a `Series`. To format this better, we need to learn about accessing individual rows in a series.

1.5 Rows and the index

In an array, you ask for data using an the location (as an int) of the item you want. You can do this in pandas using `iloc`. Add this to the end of your program:

```
# First bike
print("\n*** First Bike ***")
row = df.iloc[0]
print(f"{row}")
```

When you run it, you will see the attributes of the first row of data:

```
*** First Bike ***
brand      GT
size       57
purchase_price  277.99
purchase_date 1986-09-07
status     available
```

Name: 5636248, dtype: object

Notice that the data coming back is actually another series.

The last line says that the name (the value for the index column) for this row is 5636248. In pandas, we usually use this to locate particular rows. For example, there is a row with `bike_id` equal to 2969341. Let's ask for one entry from the

```
print("\n*** Some Bike ***")
brand = df.loc[2969341]['brand']
print(f"brand = {brand}")
```

Now, you will see the information about that bike:

```
*** Some Bike ***
brand = Cannondale
```

pandas has a few different ways of getting to that value. All of these get you the same thing:

```
brand = df.loc[2969341]['brand'] # Get row, then get value
brand = df['brand'][2969341]     # Get column, then get value
brand = df.loc[2969341, 'brand'] # One call with both row and value
```

1.6 Changing data

One of your attributes needs cleaning up. Every bike should have a status, and it should be one of the following strings: "available", "rented", or "broken". Get counts for each unique value in status:

```
print("\n*** Status ***")
series = df["status"]
missing = series.isnull()
print(f"{missing.sum()} bikes have no status.")
series_counts = series.value_counts()
for value in series_counts.index:
    print(f"{series_counts.loc[value]} bikes are \"{value}\"")
```

This will show you:

```
*** Status ***
```

```
7 bikes have no status.  
389 bikes are "rented"  
304 bikes are "broken"  
296 bikes are "available"  
1 bikes are "Flat tire"  
1 bikes are "Available"
```

Right away, we can see two easily fixable problems: Someone typed “Available” instead of “available”. Right after you read the CSV in, fix this in the data frame:

```
mask = df['status'] == 'Available'  
print(f"{mask}")  
df.loc[mask, 'status'] = 'available'
```

When you run this, you will see that the mask is a series with `bike_id` as the index and `False` or `True` as the value, depending on whether the row’s status was equal to “Available”.

When you use `loc` with this sort of mask, you are saying “Give me all the rows for which the mask is `True`.” So, the assignment only happens in the one problematic row.

Let’s get rid of the mask variable and do the same for turning `Flat tire` into `Broken`:

```
df.loc[df['status'] == 'Available', 'status'] = 'available'  
df.loc[df['status'] == 'Flat tire', 'status'] = 'broken'
```

Now those problems are gone:

```
7 bikes have no status.  
389 bikes are "rented"  
305 bikes are "broken"  
297 bikes are "available"
```

What about the rows with no values for status? If we were pretty certain that the bikes were available, we could just set them to ‘available’:

```
missing_mask = df['status'].isnull()  
df.loc[missing_mask, 'status'] = 'available'
```

Or maybe we would print out the IDs of the bikes so that we could go look for them:

```
missing_mask = df['status'].isnull()
```

```
missing_ids = list(df[missing_mask].index)
print(f"These bikes have no status:{missing_ids}")
```

However, let's just keep the rows where the status is not null:

```
missing_mask = df['status'].isnull()
df = df[~missing_mask]
```

At the end of your program, write out the improved CSV:

```
df.to_csv('bikes2.csv')
```

Run the program and open `bikes2.csv` in a text editor.

1.7 Derived columns

Let's say that you want to add a column with age of the bicycle in days:

```
bike_id,brand,size,purchase_price,purchase_date,status,age_in_days
5636248,GT,57,277.99,1986-09-07,available,13061
4156134,Giant,56,201.52,2005-01-09,rented,6362
7971254,Cannondale,54,292.25,1978-02-28,available,16174
```

Your first problem is that the `purchase_date` column looks like a date, but really it is a string. So, you need to convert it to a date. You can do this by applying a function to every item in the series:

```
df['purchase_date'] = df['purchase_date'].apply(lambda s: datetime.date.fromisoformat(s))
```

(With pandas, there is often more than one way to do things. pandas has a `to_datetime` function that converts every entry in a sequence to a datetime object. Here is another way to convert the string column in to a date column:

```
df['purchase_date'] = pd.to_datetime(df['purchase_date']).dt.date
```

You can look up `dt` and `date` if you are curious.)

Now, we can use the same trick to create a new column with the age in days:

```
today = datetime.date.today()
df['age_in_days'] = df['purchase_date'].apply(lambda d: (today - d).days)
```

When you run this, the new `bikes.csv` will have an `age_by_date` column.

Data tables in SQL

Most organizations keep their data as tables inside a relational database management system (compared to pandas, CSV, or spreadsheets). Developers talk to those systems using a language called SQL (“Structured Query Language”).

Some relational database managers are pricey products you may have heard of before, such as Oracle or Microsoft SQL Server. Some are free, such as PostgreSQL or MySQL. These are server software that client programs talk to over a company’s network.¹

There is a library, called `sqlite`, that lets us create files that hold tables. We can use SQL to create, edit, and browse those tables. `sqlite` is free, fast, and very easy to install. We will use `sqlite` instead of a networked database management system.

If you look in your digital resources, you will find a file called `bikes.db`. We created this file using `sqlite`, and now you will use `sqlite` to access it.

In the terminal, get to the directory where `bikes.db` lives. To open the `sqlite` tool on that file:

```
> sqlite3 bikes.db
```

(If your system complains that there is no `sqlite3` tool, you need to install `sqlite`. See this website: <https://sqlite.org/>)

Please follow along: type each command shown here into the terminal and see what happens.

We mostly run SQL commands in this tool, but there are a few non-SQL commands that all start with a period. To see the tables and their columns, you can run `.schema`:

```
sqlite> .schema
CREATE TABLE bike (bike_id int PRIMARY KEY, brand text, size int,
                    purchase_price real, purchase_date date, status text);
```

¹It should be noted that many notetaking and file storage applications allow you to run SQL-like queries to search your files to meet certain criteria. For example, Obsidian, the Markdown notetaking app, has a plugin called Dataview, which allows you to run searches for notes matching certain *metadata* attribute criteria. See more about this here: [FIXME obsidian dv link](#)

That is the SQL command that we used to create the bike table. You can see all the columns and their types.

You want to see all the rows of data in that table?

```
sqlite> select * from bike;
4997391|GT|57|269.61|2009-05-03|rented
5429447|Cannondale|50|215.91|2002-02-17|broken
5019171|Trek|58|251.17|1985-07-11|rented
3000288|Cannondale|57|211.08|1993-01-05|broken
880965|GT|52|281.75|1995-08-02|available
...
```

You will see 1000 rows of data!

The SQL language is not case-sensitive, so you can also write it like this:

```
sqlite> SELECT * FROM BIKE;
```

Often, you will see SQL with just the SQL keywords in all caps:

```
sqlite> SELECT * FROM bike;
```

The semicolon is not part of SQL, but it tells sqlite that you are done writing a command and that it should be executed.

SQL lets you choose which columns you would like to see. The asterisk (FIXME) used above signifies all columns, and the `bike_id`, `brand` only gets the bike's id and brand from the dataframe:

```
sqlite> SELECT bike_id, brand FROM bike;
4997391|GT
5429447|Cannondale
5019171|Trek
3000288|Cannondale
...
```

Using `WHERE`, SQL lets you use conditions to decide which rows you would like to see, and can be combined with the common operators `AND`, `OR`, and `NOT`:

```
sqlite> SELECT * FROM bike WHERE purchase_date > '2009-01-01' AND brand = 'GT';
4997391|GT|57|269.61|2009-05-03|rented
326774|GT|56|165.0|2009-06-27|available
264933|GT|52|302.43|2009-07-09|available
```

```
5931243|GT|55|173.56|2009-11-26|rented
4819848|GT|51|221.71|2009-12-11|rented
9347713|GT|52|232.32|2009-06-13|available
3019205|GT|58|262.94|2009-08-22|available
```

Using DISTINCT, SQL lets you get just one copy of each value:

```
sqlite> SELECT DISTINCT status FROM bike;
rented
broken
available
```

```
Busted
Flat tire
good
out
Rented
```

You can also edit these rows. For example, if you wanted every status that is Busted to be changed to broken, you can use an UPDATE statement with a SET:

```
sqlite> UPDATE bike SET status='broken' WHERE status='Busted';
sqlite> SELECT DISTINCT status FROM bike;
rented
broken
available
Flat tire
good
out
Rented
```

You can insert new rows:

```
sqlite> INSERT INTO bike (bike_id, brand, size, purchase_price, purchase_date, status)
...> VALUES (1, 'GT', 53, 123.45, '2020-11-13', 'available');
sqlite> SELECT * FROM bike WHERE bike_id = 1;
1|GT|53|123.45|2020-11-13|available
```

Note that the `bike_id` here must be *unique*.

You can delete rows:

```
sqlite> DELETE FROM bike WHERE bike_id = 1;
sqlite> SELECT * FROM bike WHERE bike_id = 1;
```

To get out of sqlite, type `.exit`.

Exercise 1 SQL Query

Execute an SQL query that returns the `bike_id` (no other columns) of every Trek bike that cost more than \$300.

Working Space

Answer on Page 25

2.1 Using SQL from Python

The people behind sqlite created a library for Python that lets you execute SQL and fetch the results from inside a python program.

Let's create a simple program that fetches and displays the bike ID and purchase date of every Trek bike that cost more than \$300.

Create a file called `report.py`:

```
import sqlite3 as db

con = db.connect('bikes.db')
cur = con.cursor()

cur.execute("SELECT bike_id, purchase_date FROM bike WHERE purchase_price > 330 AND brand='Trek'")
rows = cur.fetchall()

today = datetime.date.today()
for row in rows:
    print(f"Bike {row[0]}, purchased {row[1]}")

con.close()
```

When you execute it, you should see:

```
> python3 report.py
Bike 4128046, purchased 2007-08-06
Bike 7117808, purchased 1995-03-12
Bike 7176903, purchased 1986-07-03
```

Bike 827899, purchased 2009-03-14
Bike 363983, purchased 1970-08-16

Diodes, Rectifiers, and Photovoltaics

Early in this course, we discussed some basic electronics components: resistors, capacitors, and inductors. In this chapter, we are going to introduce a semiconductor: *Diodes* ensure that current flows in only one direction.

You will learn a lot about semiconductors from studying diodes. We can then apply what you have learned to solar cells, transistors, and integrated circuits like CPUs.

We are going to focus on semiconductors made of silicon. You can make semiconductors out of other materials, but silicon has proved to be very versatile and cost-effective.

3.1 Silicon

A silicon atom has 14 protons and 14 electrons. The inner 10 electrons are arranged in two complete and stable shells. The outer four electrons are in an incomplete shell, which allows them to form bonds with other atoms. We call these *valence electrons* because they aren't tightly bound to their atoms. A complete outer shell would have eight electrons.

Atoms can share electrons to get complete shells. This is called *covalent bonding*. When atoms share electrons, they form a *covalent bond*. The electrons are shared between the atoms, and the atoms are held together by the shared electrons.

In a silicon crystal, each silicon atom shares its outer four electrons with four neighboring atoms. With its four electrons and an electron borrowed from each of its four neighbors, each silicon atom has eight electrons in its outer shell. This is a complete shell, and the atoms are stable. This is why glass is hard and stable. It also why electricity doesn't flow through glass – all the electrons are tightly bound to their atoms.

Phosphorus has five valence electrons. If you put a phosphorus atom into silicon crystal, it shares four of its electrons with its neighbors, but it has one electron left over after the complete shell of eight is built. As you add more phosphorus to a silicon crystal, it conducts more electricity. We call this *doping*.

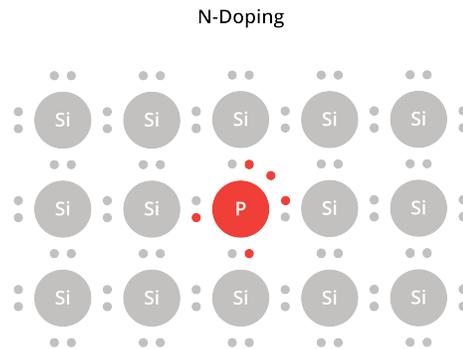


Figure 3.1: Phosphorus atom in a silicon crystal

Because this kind of doping frees up some electrons, which are negatively charged, we call this *n-doping*.

Boron has only three valence electrons. If you put a boron atom into silicon crystal, it shares three of its electrons with its neighbors, which leaves a "hole" that would like to be filled by an electron. We call this *p-doping*.

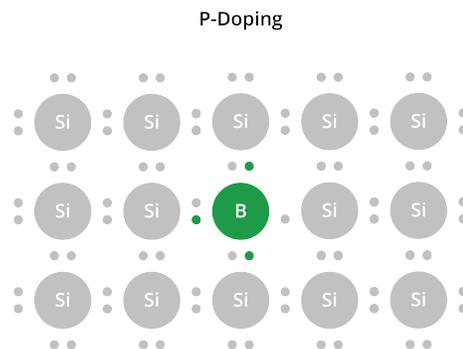


Figure 3.2: Boron atom in a silicon crystal

Both n-doped and p-doped silicon are conductors of electricity. In n-doped silicon, we say the loose electrons are the *carriers* of the current. In p-doped silicon, we say the holes are the carriers.

Note that doped silicon is still mostly silicon. In normal doping, maybe 1 in 10,000 atoms are replaced with boron or phosphorus.

3.2 Diodes

A diode is a semiconductor device that allows current to flow in one direction only.

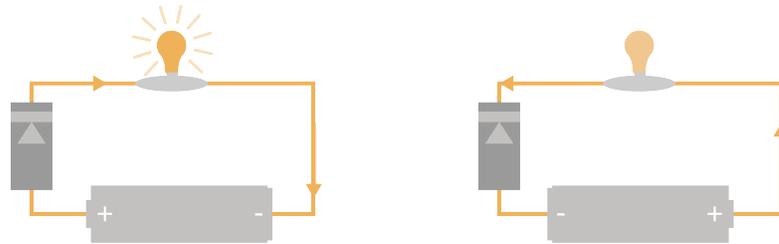


Figure 3.3: Current flows in one direction through the diode, not in the other direction.

Diodes can be made with silicon that has been n-doped and p-doped. A diode looks like this:



Figure 3.4: What Diodes Look Like

We say that current flows from the positive terminal to the negative terminal of a battery. However, remember that the electrons are actually moving from the negative terminal to the positive terminal. This is a common source of confusion in any discussion of diodes.

One end of the diode is called the *anode*. The anode is p-doped. The other end is called the *cathode*. The cathode is n-doped.



Figure 3.5: A p-n Diode

The region close to where the anode and cathode meet is called the *depletion region*. In this space, the spare electrons in the cathode (the n-doped material) migrate over to fill the holes in anode (the p-doped material). This creates an electric field: there are more electrons than protons on the anode side, and more protons than electrons on the cathode side. (How strong is this field? It is typically about 0.7 volts.)

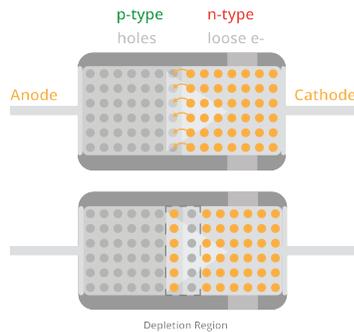


Figure 3.6: The depletion region in a diode

What happens if we overpower that 0.7 volts? What if we put 2 volts in the opposite direction? This shrinks the depletion region and the diode becomes a conductor. For example, if the negative end of a battery is connected to the cathode of a diode and the positive end is connected to the anode, current will flow through the diode.

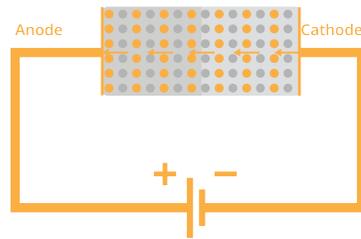


Figure 3.7: Electrons flow through the diode

If we reverse the battery, the depletion region expands, and the diode becomes an insulator.

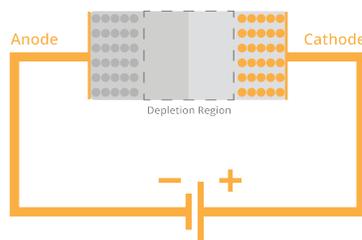


Figure 3.8: Reverse the electric field? No flow.

Here is the symbol for a diode:

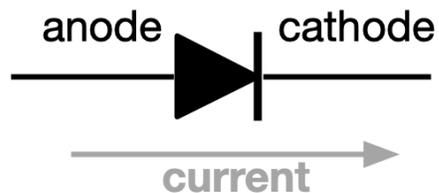


Figure 3.9: Diode Symbol

Current will flow in the direction of the arrow. It will not flow in the opposite direction.

3.2.1 Rectifiers

One of the really useful things we can do with a diode is to convert alternating current (AC) into direct current (DC). This is called rectification. Remember that AC pushes and pulls the electrons back and forth in the wire. Rectification ensures that the electrons only move in one direction.

A *half-wave rectifier* uses a single diode to ensure that the current flows in one direction through the load.



Figure 3.10: Half-wave Rectifier

Current flows in the direction of the arrow on the diode. If the voltage is reversed, the diode will block the current from moving.

A *full-wave rectifier* uses four diodes to swap the direction of the current through the load when necessary.

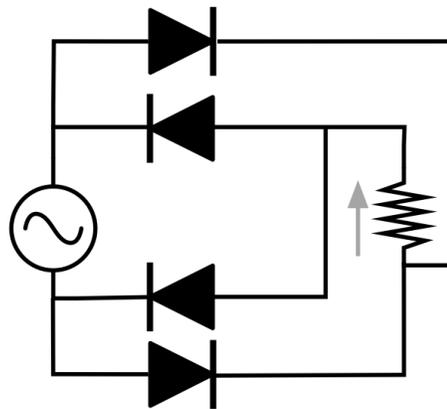


Figure 3.11: Full-wave Rectifier

3.2.2 Photovoltaics

Photovoltaics (or "Solar cells") are similar to diodes. They have a p-n junction with a 0.7 volt electrical field. The difference is that sunlight is allowed to shine on the cathode

side. When an atom is struck by a photon of light energy, the electron becomes excited. Sometimes it is excited enough to leap across that 0.7 volt barrier.

When it leaps across, there is an extra electron on the anode side of the barrier and an extra proton on the cathode side. The electron is pushed back to where it came from, but it can't go back over the barrier (unless more than 0.7 volts of charge have built up across the barrier). So, it goes the long way: through the load.

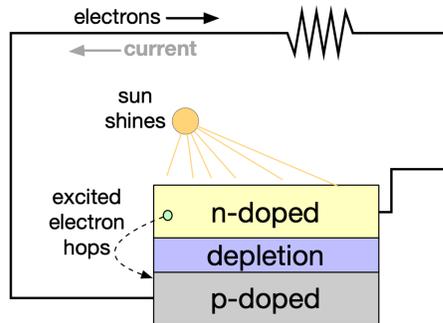


Figure 3.12: Solar Cell

Answers to Exercises

Answer to Exercise 1 (on page 14)

```
SELECT bike_id FROM bike WHERE purchase_price > 330 AND brand='Trek'
```




INDEX

anode, [19](#)

boron, [18](#)

cathode, [19](#)

datatypes, [3](#)

depletion region, [20](#)

diode, [17](#)

doping, [17](#)

full-wave rectifier, [22](#)

half-wave rectifier, [22](#)

n-doped, [18](#)

null, [3](#)

p-doped, [18](#)

pandas, [4](#)

phosphorus, [17](#)

photovoltaics, [23](#)

rectifier, [22](#)

silicon, [17](#)

solar cells, [23](#)

SQL, [11](#)

SQLite, [11](#)