

ExAx: The Accelerator

User's Reference

2025-11-07, **draft**



— Fast and Reproducible Data Processing —

NOTE: Parts of this document are out of date.
We are currently prioritising the online documentation available at

<https://exax.readthedocs.io>

Anders Berkeman, Carl Drougge, and Sofia Hörberg

version: 6c988078

Contents

1	Introduction	8
1.1	Main Design Goals	9
2	Overview	10
2.1	High Level View	11
2.2	Interacting with the Accelerator	11
2.3	Jobs	11
2.3.1	A “Hello, World” Example	12
2.3.2	Jobs Can Only be Run Once	12
2.3.3	A Quick Look Under the Hood	13
2.3.4	Workdirs and Sharing Jobs	13
2.3.5	Linking Jobs	13
2.3.6	Job Input Parameters	14
2.3.7	Parallel Processing	14
2.3.8	Job Execution Flow	15
2.4	Datasets: Storing Data	16
2.4.1	Importing Existing Data from a File	16
2.4.2	Linking Datasets, Chaining	17
2.4.3	Adding New Columns to a Dataset	17
2.4.4	Multiple Datasets in a Job	18
2.4.5	Parallel Dataset Access and Hash Partitioning	18
2.4.6	Dataset Column Types	18
2.4.7	Dataset Attributes	19
2.5	Iterators: Working with Data	19
2.5.1	Iterator Basics	19
2.5.2	Parallel Execution	19
2.5.3	Iterating over Several Columns	20
2.5.4	Iterating over Dataset Chains	20
2.5.5	Job Execution Flow and Result Passing	21
2.5.6	Job Parameters	22
2.6	A Class Based Programming Model	22
3	Basic Build Scripting	23
3.1	Build Scripts	24
3.1.1	Building a Job: <code>urd.build()</code>	24
3.1.2	Connecting Jobs	25
3.1.3	Replaying Build Scripts	25
3.2	Working with Build History: <code>urd.joblist</code>	25
3.2.1	Printing a <code>JobList</code> : <code>urd.joblist.pretty</code>	25
3.2.2	Finding Jobs in a <code>Joblist</code>	26
3.2.3	Indexing and Slicing a <code>JobList</code>	26
3.3	Configuration Information: <code>urd.info</code>	26
3.4	Summary	27

4	Jobs	28
4.1	Definitions	29
4.1.1	Methods and Jobs	29
4.1.2	Jobids	29
4.1.3	Work Directories and Job Directories	29
4.1.4	The <code>Job</code> and <code>CurrentJob</code> Convenience Wrappers	30
4.2	Python Packages	30
4.2.1	A New Package using <code>auto-discover</code>	30
4.2.2	A New Package using <code>methods.conf</code>	31
4.3	Method Source Files	31
4.3.1	Creating a New Method	31
4.3.2	Making Methods Executable: <code>methods.conf</code>	32
4.4	Job Building or Job Recycling	32
4.4.1	Job Already Built Check	32
4.4.2	Depend on More Files: <code>depend_extra</code>	33
4.4.3	Avoiding Rebuild: <code>equivalent_hashes</code>	33
4.4.4	Forcing a Rebuild	33
4.5	Invoking the Python Debugger	34
4.6	Method Execution	34
4.6.1	Input Parameters	34
4.6.2	Functions Reserved for Execution Flow	34
4.6.3	Execution Order	34
4.6.4	Function Arguments	35
4.6.5	Parallel Processing: The <code>analysis()</code> Function and Slices	35
4.6.6	Return Values	35
4.6.7	Automatically Merging Results from <code>analysis()</code>	35
4.6.8	Standard Out and Standard Error	36
4.7	The <code>Job</code> and <code>CurrentJob</code> Classes	36
4.7.1	Writing and Reading Serialised Data	36
4.7.2	Writing and Reading Serialised Data in Parallel	37
4.7.3	General File Access	37
4.7.4	Accessing a Job's Return Value	37
4.7.5	Accessing a Job's Datasets	38
4.7.6	Accessing a Job's Options and Parameters	38
4.7.7	Accessing Job Output	39
4.7.8	Reading Post Data	39
4.8	From Jobs to Datasets and Back	40
4.8.1	From Dataset to Job	40
4.8.2	From Job to Dataset	40
4.8.3	From Dataset to Dataset (in same Job)	40
4.9	Method Input Parameters	40
4.9.1	Input Jobs	40
4.9.2	Input Datasets	41
4.9.3	Input Options	41
4.9.4	A Specific File From Another Job: <code>JobWithFile</code>	41
4.10	Subjobs	42
4.11	Formal Option Rules	42
4.11.1	Options with no Type	43
4.11.2	Scalar Options	43
4.11.3	String Options	44
4.11.4	Enumerated Options	44
4.11.5	List and Set Options	44
4.11.6	Date and Time Options	45
4.11.7	More Complex Stuff: Types Containing Types	45
4.12	Jobs - a Summary	45

5	Datasets	47
5.1	Dataset Internals	48
5.2	Slicing and Hashing	49
5.3	Chaining	49
5.4	Dataset as Job Input Parameter	49
5.5	Datasets from Job Objects	50
5.6	Dataset Properties	50
5.6.1	Dataset Name	50
5.6.2	Column Names	50
5.6.3	Column Properties	51
5.6.4	Rows per Slice	51
5.6.5	Dataset Shape	51
5.6.6	Hashlabel	51
5.6.7	Filename and Caption	51
5.7	Operations on Chains	52
5.8	Column Data Types	52
5.8.1	Arbitrary precision numbers: <code>number</code>	53
5.8.2	Standard Fixed Size Numbers	53
5.8.3	Booleans	54
5.8.4	Types Relating to Time	54
5.8.5	String Types	54
5.8.6	Raw Data	54
5.8.7	Complex Numbers	54
5.8.8	JSON	54
5.8.9	Pickle	54
5.8.10	<code>None</code> -Handling	54
5.8.11	Does Hash Partitioning Depend on the Data Type?	55
5.9	Creating a New Dataset	55
5.9.1	Create in <code>prepare()</code> + <code>analysis()</code>	55
5.9.2	Create in <code>synthesis()</code>	56
5.9.3	Creating a Dataset Chain in one Job	56
5.9.4	Completing Dataset Creation	56
5.9.5	Datasets Created by Subjobs	57
5.9.6	Creating Hash Partitioned Datasets	57
5.9.7	Column Name Restrictions	57
5.9.8	More Advanced Dataset Creation	58
5.10	Appending New Columns to an Existing Dataset	58
5.10.1	Appending New Columns in Analysis	58
5.10.2	Appending New Columns in Synthesis	58
5.11	Standard Methods on Datasets	59
6	Data Iterators	61
6.1	The Three Iterators	62
6.1.1	Iterator Arguments	63
6.2	Basic Iteration	64
6.2.1	Parallel Iterator Invocation	64
6.2.2	Sequential Iterator Invocation	64
6.2.3	Special Case, Round Robin Iteration	64
6.2.4	Special Cases, Iterating Over All or a Single Column	64
6.2.5	Iterate Over Chains	65
6.3	Halting Iteration	65
6.3.1	Halting Using <code>length</code>	65
6.3.2	Halting Using <code>stop_ds</code>	65
6.3.3	Halting Using Another Job's Input Parameters	65
6.4	Iterating Over a Data Range	66
6.5	Iterating in the Reverse Direction	66
6.6	Hash Partitioned Datasets and on-the-fly Partitioning	66
6.7	Callbacks	67
6.7.1	Skipping Datasets and Slices from Callbacks	67

7	High Level Control: Urd	69
7.1	Introduction to Urd	70
7.2	A Simple Use Case	70
7.2.1	The Import Build Script	70
7.2.2	The Analysis Build Script	70
7.2.3	Urd Provides Separation	70
7.2.4	Transparency	71
7.3	Local or Shared Urd Server	71
7.4	Urd Sessions and Lists	71
7.5	A First Urd Query	71
7.6	The Contents of a Stored Session	72
7.7	Urd Sessions: <code>.begin()</code> and <code>.finish()</code>	73
7.7.1	What if a Build Script is Run Again?	73
7.7.2	Timestamp Definition and Resolution	73
7.8	Retrieving an Urd Session	74
7.8.1	Finding an Exact or Closest Match: <code>.get()</code>	74
7.8.2	Finding the Latest Session: <code>.latest()</code>	75
7.8.3	Finding the first item: <code>.first()</code>	75
7.9	Retrieving an Urd Session without Recording	75
7.10	Aborting an Urd Session: <code>.abort()</code>	75
7.11	Truncating and Updating	76
7.11.1	Updating the Last Item	76
7.11.2	Truncating a List	76
7.11.3	Truncation Consequences: Ghosts	76
7.12	More Search Functions	76
7.12.1	Listing all Urd Lists: <code>.list()</code>	76
7.12.2	Listing all Items After a Specific Timestamp: <code>.since()</code>	77
7.13	Building Jobs: <code>.build()</code>	77
7.13.1	Building Chained Jobs: <code>urd.build_chained()</code>	78
7.14	Changing workdir: <code>.set_workdir()</code>	79
7.15	Profiling a Build Script: <code>.print_exeetimes()</code>	79
7.16	Passing Flags from the Command Line	79
7.17	The <code>ax urd</code> Command and the Board	80
7.18	Multi-User Urd Consistency	80
8	Standard Methods	81
8.1	<code>csvimport</code> – Importing Data Files	82
8.1.1	Options	82
8.1.2	Datasets	83
8.1.3	Bad Lines	83
8.1.4	Output	83
8.1.5	Line Numbers	83
8.1.6	Limitations	84
8.1.7	Example Invocation	84
8.2	<code>csvimport_zip</code> – Importing zip Archives	85
8.2.1	Options	85
8.2.2	Example invocation	85
8.3	<code>dataset_type</code> – Typing Datasets	87
8.3.1	Datasets	87
8.3.2	Options	87
8.3.3	Example Invocation	88
8.3.4	Typing	88
8.4	<code>csvexport</code> – Exporting Text Files	91
8.4.1	Example Invocation	91
8.5	<code>dataset_hashpart</code> – Hash Partition a Dataset	93
8.5.1	Example Invocation	93
8.5.2	Hashing Details	93
8.5.3	Notes on Chains	93
8.6	<code>dataset_filter_columns</code> – Removing Columns from a Dataset	94

8.7	<code>dataset_sort</code> – Sorting a Dataset	95
8.7.1	Sorting <code>None</code> and <code>NaN</code> values	95
8.7.2	Spreading of Left Over Values	95
8.7.3	The Trigger Column	95
8.7.4	A Practical Limitation	96
8.8	<code>dataset_checksum</code> , <code>dataset_checksum_chain</code>	97
8.9	<code>dataset_merge</code> – Merge Several Datasets into One	98
8.10	<code>dataset_unroundrobin</code> – Transpose a Dataset	99
9	Accelerator Command Line Tools	100
9.1	Quick Summary of Commands	101
9.2	Job and Dataset References for Command Line Tools	102
9.3	Initialisation: <code>ax init</code>	103
9.4	Accelerator Server: <code>ax server</code>	104
9.5	Running Build Scripts: <code>ax run</code>	104
9.6	Dataset Information: <code>ax ds</code>	105
9.7	Look at Specific Data in a Dataset: <code>ax grep</code>	107
9.8	Information about Methods: <code>ax method</code>	108
9.9	List all build scripts: <code>ax script</code>	109
9.10	Information about Jobs: <code>ax job</code>	109
9.11	Aliases	109
9.12	Show Dataset Contents: <code>ax cat</code>	109
9.13	The Urd Job Database Server: <code>ax urd-server</code>	109
9.13.1	Authorization to Urd	110
9.14	Aborting a Running Job: <code>ax abort</code>	110
9.15	Current Version: <code>ax version</code>	110
9.16	Workdir Information: <code>ax workdir</code>	110
9.17	The Board Web Server: <code>ax board-server</code>	110
9.18	Displaying Urd Database Content: <code>ax urd</code>	111
A	Miscellaneous	112
A.1	Workdirs	113
A.1.1	Creating a Workdir	113
A.2	Status (Progress) Reporting	113
A.3	Generating Progress Messages: the <code>status</code> Module	113
A.4	Working with Relative Paths	114
A.4.1	The <code>input_directory</code>	114
A.4.2	The <code>result_directory</code>	114
A.5	Concurrency	115
B	Setup and Installation	116
B.1	Install the Accelerator	117
B.1.1	Using the <code>pip</code> command	117
B.2	Set up a New Project	117
B.3	Run the Tests	117
B.4	Server Configuration File	118
B.5	Setting up a Standalone Urd-server	120
B.5.1	Starting Urd	120
B.5.2	The <code>passwd</code> file	120
B.5.3	Connecting to a Remote Urd Server	121
B.5.4	The Urd Database	121
B.6	Setting up a Standalone Board-server	121
C	Classes	122
C.1	The Job and <code>CurrentJob</code> Classes	123
C.1.1	<code>Job.chain()</code>	124
C.1.2	<code>Job.dataset()</code>	124
C.1.3	<code>Job.filename()</code>	124
C.1.4	<code>Job.files()</code>	124
C.1.5	<code>Job.json_load()</code>	124

C.1.6	Job.link_result()	124
C.1.7	Job.load()	125
C.1.8	Job.open()	125
C.1.9	Job.output()	125
C.1.10	Job.withfile()	126
C.1.11	CurrentJob.input_filename()	126
C.1.12	CurrentJob.json_save()	126
C.1.13	CurrentJob.open_input()	126
C.1.14	CurrentJob.register_file()	127
C.1.15	CurrentJob.save()	127
C.1.16	Sliced Files	127
C.1.17	File Persistence	127
C.2	The JobWithFile Class	128
C.3	The JobList Class	129
C.3.1	JobList.find()	129
C.3.2	JobList.get()	129
C.3.3	JobList.print_execetimes()	129
C.4	The Dataset Class	130
C.4.1	Dataset.min(), Dataset.max()	130
C.4.2	Dataset.link_to_here()	130
C.4.3	Dataset.merge()	131
C.4.4	Dataset.chain()	131
C.5	The DatasetChain Class	132
C.5.1	DatasetChain.min(), DatasetChain.max()	132
C.5.2	DatasetChain.lines()	132
C.5.3	DatasetChain.column_counts()	132
C.5.4	DatasetChain.column_count()	132
C.5.5	DatasetChain.with_column()	133
C.6	The DatasetWriter Class	134
C.6.1	DatasetWriter.add()	134
C.6.2	DatasetWriter.hashcheck()	134
C.6.3	DatasetWriter.set_slice()	134
C.6.4	DatasetWriter.enable_hash_discard()	135
C.7	The Urd Class	136
C.7.1	Urd.get()	136
C.7.2	Urd.latest()	136
C.7.3	Urd.first()	136
C.7.4	Urd.peek()	137
C.7.5	Urd.peek_latest()	137
C.7.6	Urd.peek_first()	137
C.7.7	Urd.since()	137
C.7.8	Urd.list()	137
C.7.9	Urd.begin()	137
C.7.10	Urd.abort()	138
C.7.11	Urd.finish()	138
C.7.12	Urd.truncate()	138
C.7.13	Urd.set_workdir()	138
C.7.14	Urd.build()	138
C.7.15	Urd.build_chained()	138
C.7.16	Urd.warn()	139

Chapter 1

Introduction

DRAFT

ExAx, or the Accelerator, is a tool for fast reproducible data processing, capable of working at high speed with terabytes of data with billions of rows on a single computer. The speed in combination with its unique capabilities to ensure reproducibility makes the Accelerator a good choice for tasks where it is important to keep track of how data, code, and results are connected. Furthermore, using the Accelerator saves time, since results are *re-used*, and not re-computed, whenever possible.

Typical applications include all kinds of data analysis work as well as live production systems for tasks such as recommender systems, and more. The Accelerator has a small footprint, few dependencies, and runs on laptops as well as rack servers.

The Accelerator was first used in 2012, and has been continuously developed and improved since. It has been in use in projects for companies like *Safeway*, *Starbucks*, *eBay*, *Ericsson*, and *Vodafone*, among others. Most projects have been related to data analysis, some to optimisation, and some projects have been recommender systems running live for years. The Accelerator has been the core of these projects. In 2016, the Accelerator was acquired by Ebay, who contributed it to the open source community early 2018.

The sizes of typical data sets used in these projects range from a few hundred lines up to several tens of billions rows and many columns. The number of items in a dataset used in a live system was well above 10^{11} , and this was handled with ease on a *single* 32 core computer.

More than 1600 commits have been removed to clean up the open version of the code base, and about 1600 new commits have been created since the Accelerator was open sourced. The Accelerator is written in Python, with the exception of critical parts that are written in the C programming language.

The authors are Anders Berkeman, Carl Drougge, and Sofia Hörberg. Extensive testing has been done by Stefan Weiland, and Andreas Ottosson has been proofreading.

1.1 Main Design Goals

The Accelerator was originally developed for reproducible, transparent, and maximally fast data processing using all available CPU cores. It is designed bottom up for high performance and simplicity, and the main design goals are:

Parallel processing should be made simple. Modern computers come with several CPU cores. It should be straightforward to use them to speed up computations.

Data rates should be as high as possible, i.e. close to the hardware bounds. It should be possible to process *large datasets on inexpensive hardware*, including laptops.

All processing steps should be *reproducible*. The Accelerator maps any generated result to its corresponding input data and processing source code.

Never recompute old results, always “recycle” old jobs, when possible. *Sharing* computed results between multiple users should be effortless.

Organise and keep track of all jobs, input files, intermediate files, and results to make it easy to work with projects having 100.000s of input files and lots of programs and scripts processing them.

In addition, the Accelerator is originally designed to be used at all levels of a project, including data engineering, data analysis, algorithm development, as well as production. Nevertheless, it excels as a pure data analysis or data processing tool.

Chapter 2

Overview

This chapter presents an overview of the Accelerator's features in a rather non-formal way. It is based on an article published on the eBay Tech Blog website.

DRAFT

2.1 High Level View

The Accelerator is a client-server based application, and from a high level, it can be visualised like in figure 2.1.

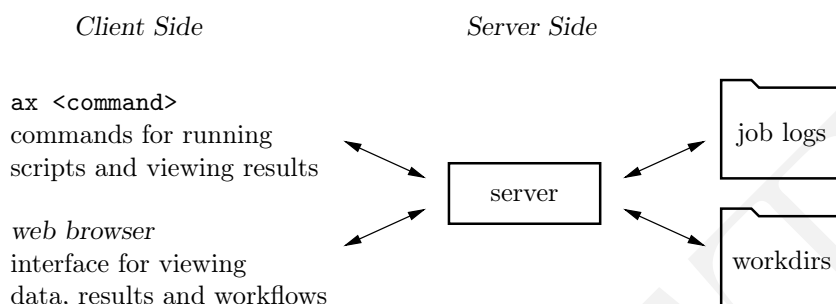


Figure 2.1: High level view of the Accelerator framework. See text for details.

The left part of the figure shows the *client* side, with shell commands and web interface. On the right side is the *server* and its databases. The Accelerator executes *build scripts*, that start *jobs* on the server. These jobs are stored in the server's job databases for later retrieval. There are two databases, one for storing everything related to a job's execution, called the *workdir*, and one transaction log used for storing meta-information about built jobs. The *workdir* will contain all inputs, source code, parameters, and outputs of all executed jobs, whereas the job log database ensures reproducibility and transparency, and it will be further discussed in chapter 7.

2.2 Interacting with the Accelerator

There are two ways to interface the Accelerator, using

- the command line, and
- using a web browser.

The command line interface is based on the `ax` shell command. It can be used for a large variety of operations, including starting execution of scripts, listing or grepping data in datasets, finding particular jobs, and a lot more.

The web interface, called *Accelerator Board*, is used to inspect jobs, datasets, and meta-information. Resulting files such as images, text files, pdfs, video files and so on are rendered in the web browser.

2.3 Jobs

A problem that is solved using the Accelerator is typically partitioned into smaller programs called *methods*. A method is a Python program with one or a few special functions used to execute code sequentially or in parallel, and to pass parameters and results.

For each method that is executed by the Accelerator, a unique new directory, called *job directory*, is created in which everything related to the execution, such as inputs, outputs, source code etc., is stored. An executed method is called a *job*. To be truly transparent, the job directory will also contain profiling information, Python interpreter version and other related information.

The Accelerator has a database that keeps track of all jobs that has been run. This avoids unnecessary re-computing in favour of reusing previously computed results. Reusing jobs does not only speed up processing and encourage incremental design, but also makes it transparent which code and which data that was used for any particular result, minimising these kind of uncertainties to zero.

2.3.1 A “Hello, World” Example

Here’s an example of a very simple method. It does not take any input parameters and does almost nothing, it will just return the string “hello world” and exit. The example is stored in a file named `a_hello_world.py`¹:

```
def synthesis():  
    return "hello world"
```

The `synthesis()`-function is reserved for code that is to be executed once without parallelisation. The method is stored in a separate file. In order to get the method to execute, it is called from a *build script* looking something like this

```
def main(urd):  
    job = urd.build('hello_world')  
    print(job.load())
```

and stored in a file `build_mytest.py`. The `main()`-function is called when the build script is run. The input parameter `urd` is an object that contains functions for job building and organisation, and is described in chapter 7 and section C.7. The build script is executed by typing a shell command like this

```
ax run mytest
```

In the build script, the `urd.build()`-call will instruct the server to execute the `hello_world` method. During its execution, a job directory will be created that contains everything associated with the build process. When the `build()`-call is completed, a job object, of type `Job`, is returned to the program. This object provides a convenient interface to the data in the corresponding job directory, and contains member functions such as `.load()`, that is used in the example to read back the returned value from the job.

Note that what is returned by the method is stored in a file on disk. There is no need to know the name of the file, since its contents can be retrieved by the job objects `load()` function.

2.3.2 Jobs Can Only be Run Once

If the build script is executed again, the `hello_world` job will *not* be re-built, because the Accelerator remembers that the job has been built in the past, and all associated information is already stored in a job directory. Instead, the Accelerator immediately returns a job object representing the previous run. This means that from a user’s perspective, there is no visible difference between running a job for the first time or re-using results from an existing run! In order to have the method execute again, either its source code or input parameters has to change. If there are such changes, the method will be re-executed, and a new job will be created that reflects these changes.

¹There’s a good reason for the file name prefix, see section ??.




Figure 2.2: A simple hello world program, represented as graph and job directory. Building the `hello_world` method creates a *job* `test-0`. This job is stored in a (job-) directory, using the same name, in the *test* *workdir*. Everything associated with the build is stored in this job directory.

2.3.3 A Quick Look Under the Hood

The Accelerator is fully transparent, and no data is “hidden” anywhere. There is no need to know how or where things are stored to use the Accelerator, but still it could be of interest to know where things ends up (and that it is straightforward to access things using simple shell commands.)

Figure 2.2 illustrates an execution of the `hello_world` method. The created job gets the *jobid* `test-0`, and parts of the corresponding job directory information is shown in green. Jobids are job identifiers, named by their corresponding *workdir* plus an integer counter value.

The most important files in a job directory are

- `setup.json`, containing job meta information;
- `result.pickle`, containing the returned data; and
- `method.tar.gz`, containing the method’s source code,

and in general the job directory could contain any number of files.

The Job object, which is returned by the `build` call, provides a convenient way to access files and data stored in the job directory. For example, the job’s return value can be loaded into a variable using the `.load()` function, all metadata can be pretty printed using the `ax job` shell command, and so on.

2.3.4 Workdirs and Sharing Jobs

A *workdir* is where jobs are stored. There can be more than one *workdir*, and different *workdirs* can be used separate jobs into different physical locations, or to separate different types of data from one another in the same project. The Accelerator can be set up to have any number of *workdirs* accessible, but typically only one is used for writing.

Entering the same *workdir* definition into two or more users’ configuration files, the *workdir* and its contents will be shared between the users. Each Accelerator server will update its knowledge about the contents of all *workdirs* before executing a build script, to make sure that the latest jobs are taken into account. In addition, the job log database, as described in chapter 7, is designed for efficient re-use and sharing of particularly interesting jobs. Thus, only one user has to execute a given method for all other users to have instant access to it. This can save a lot of time in a project.

2.3.5 Linking Jobs

Using jobs, complex tasks can be split into several smaller operations. Jobs can be connected to eachother so that a new job will use results from one or more previously executed jobs.

To continue the simple hello world example, assume for a second that the `hello_world`-job is computationally expensive, and should therefore be run as few times as possible to save time and energy. The job creates a result that will be used in further processing. This further processing is represented in this example by printing the result to standard output.




Figure 2.3: Job `test-0`, is used as input to the `test-1` job. Jobs are created from methods (i.e. source code), input parameters, and processing time using the `build()` call. Each job corresponds to a number of files stored in a per job unique (job) directory.

A new method `print_result` is created for this task, and it looks like this

```
jobs = ('hello_world_job',)

def synthesis():
    data = jobs.hello_world_job.load()
    print(data)
```

The “`jobs=`”-line at the top specifies that the method expects a reference to a job to be provided as input at run time. The input is arbitrarily named `hello_world_job`, and the variable `jobs.hello_world_job` will be of type `Job`, and therefore its result can be read using the `.load()` function as shown previously.

The following extended build script will run both the first and second method, and provide the first job as input to the second:

```
def main(urd):
    job1 = urd.build('hello_world')
    job2 = urd.build('print_result', hello_world_job=job1)
```

Figure 2.3 illustrates the situation. Note the direction of the arrow. Arrows illustrate pointers to previous jobs. In this example, the second job, `test-1`, has `test-0` as input parameter, i.e. `test-1` has a pointer to `test-0`. (And `test-0` does not know of any jobs run in the future, so the arrow cannot point out from it.)

The example shows how a potentially complex task can be split into several jobs, each reading intermediate results from previous ones, *without the need for explicit filenames!* The Accelerator will keep track of all job dependencies, so there is no doubt which jobs that are run when, and on which data. And since the Accelerator remembers if a job has been executed before, it will link and *re-use previous jobs when possible*, thereby saving time and energy.

Also note that a re-used job is a *proof* of that the code, input- and output data is unchanged and connected. *Using the Accelerator, there is no risk that wrong intermediate files are used by mistake*, since the files are accessed by reference to the method that created them, and not by an arbitrary user-selected filenames!

2.3.6 Job Input Parameters

In the previous section it is shown how the input parameter `jobs` is used to point to earlier jobs. There are two more types of input parameters: `options` and `datasets`, and here’s an example of how to declare all of them in a method:

```
jobs = ('hello_world_job',) # as seen in the previous example
datasets = ('transaction_log', 'access_log',)
options = dict(length=4)
```

All parameters are populated by the `build()` call in the build script. The `jobs` parameter is used to link the job to previous jobs. `options` is a dict of any kind of options the job may use, and `datasets` is, similar to `jobs`, a list of datasets. (Datasets will be introduced shortly.)

2.3.7 Parallel Processing

While the `synthesis()`-function (presented in the previous section) is used for code that is to execute exactly once, the `analysis()`-function is used for simple but powerful parallel processing. A method making use of parallel processing may look like this

```
def analysis(sliceno):
    print('This is process %d!' % (sliceno,))
```

When the method is executed, the `analysis()`-function will be forked into `slices` parallel processes (where `slices` is specified in the Accelerator’s configuration file and typically set to the number of available CPU cores of the machine). The input parameter `sliceno` is a number between zero and `slices-1` and used to identify each process.

Admittedly, this is the most basic form of parallel processing, that can only solve a subset of all kinds of parallel problems. Nevertheless, it solves common data processing problems, and it solves them in a very simple and efficient manner. The next section shows how the Accelerator makes use of the `analysis()`-function to do data processing on all CPUs in parallel. @@

2.3.8 Job Execution Flow

Execution of code in a method is either parallel or serial depending on which function is used to encapsulate it. In addition to `analysis()` and `synthesis()` that have already been introduced, there is a third one named `prepare()`. All three may exist in the same method, and at least one is required. When the method executes, they are called one after the other.

Results can be transferred from one function to the next using return values. There is also an easy way to merge results from parallel processing (introduced in section ??).

`prepare()` is executed first. The returned value is available in the variable `prepare_res`.

`analysis()` is run in `slices` parallel processes, forked after completion of `prepare()`. Common input parameters are `sliceno`, holding the number of the current process instance, and `prepare_res`. The return value for *all* processes becomes available in the `analysis_res` generator variable.

`synthesis()` is called after the last `analysis()`-process is completed. It is typically used to aggregate parallel results created by `analysis()` and takes both `prepare_res` and `analysis_res` as optional input parameters.

The return value from `synthesis()` is the default return value of the job itself. Figure 2.9 shows the execution order from top to bottom, and the data passed between functions in coloured branches.




Figure 2.4: Execution flow and result propagation in a method. `prepare()` is executed first, and its return value is available to both the `analysis()` and `synthesis()` functions. There are `slices` (a configurable parameter) number of parallel `analysis()` processes, and their outputs are available to the `synthesis()` function, which is the last to execute.

2.4 Datasets: Storing Data

The `dataset` is the Accelerator’s default storage structure for small or large quantities of data, designed for parallel processing and high performance. Datasets are built on top of jobs, so *datasets are created by methods and stored in job directories*. To ensure data integrity, datasets are *immutable*, i.e. their contents cannot be modified.

Internally, data in a dataset is stored in a row-column format. It is typically *sliced* into a fixed number of slices to allow efficient parallel access, both for reading and writing, see figure 2.5. The different columns in a dataset are stored and accessed independently, so there is no overhead in reading just a single or a smaller set of all columns of a dataset. Only the data relevant for a task is read from disk.




Figure 2.5: A dataset containing three columns, A , B , and C , stored using two slices, i.e. `slices=2`. Subscripts corresponds to line numbers. In this example, slice 0 stores the even rows, and slice 1 stores the odd rows. Each dotted box corresponds to a file, so columns could be read independently, and since there are two files for each column, any column could be processed in parallel using two independent processes.

Data in a dataset is typically accessed using a Python *iterator* that reads and streams a dataset slice from disk or memory to a CPU core for processing. Here’s a simplified example showing how to create a count the frequency of each movie in the `movie` column of a dataset. More and complete examples will be presented later in this chapter.

```
from collections import Counter
c = Counter()
for movie in dataset.iterate(None, 'movie'):
    C[movie] += 1
```

The streaming iterator approach has several benefits. It enables working on unlimited sizes of data, since only a fraction of the data is in memory at any given time. Big data analysis can be carried out with ease on a laptop. Also, iterators can be really fast, and the corresponding linear accessing of data on disk is the fastest way to access harddrive data, since it avoids time wasting seek calls.

Furthermore, datasets may be *hash partitioned*, so that slice-membership is based on the output of a hash function applied to the values of a given column. Hash partitioning is a very fast linear time operation that is used to separate data so that parallel computations can be carried out *independently* in each slice, without the need for complicated merging operations. This is another reason for the Accelerators high performance, and it is further explained in section 5.2.

The parallel processing capabilities of the Accelerator makes it possible to dispatch a set of parallel iterators, one for each slice and CPU core, in order to maximise use of all available hardware resources.

This section shows how iterators are used for reading data, how to take advantage of slicing to have parallel processing, and how to efficiently create new datasets.

For further information on datasets, visit Chapter 5.

2.4.1 Importing Existing Data from a File

A project typically starts with *importing* some data from a file on disk. The bundled method `csvimport` is designed to parse a plethora of “comma separated values”-file formats and store

the data as a *dataset*. The method takes several input options in addition to the mandatory filename to control the import process, see section ?? Here is an example invocation

```
def main(urd):
    jid = urd.build('csvimport', filename='file0.txt')
```

When executed, the created dataset will be stored in the resulting job directory, and the name of the dataset will by default be the jobid plus the string `default`. For example, if the `csvimport` jobid is `imp-0`, the dataset will be referenced by `imp-0/default`. In this case, the common case where a job only contains one dataset, the dataset could also be referenced by the job name `imp-0` for simplicity. But in general a job can contain any number of datasets and then a unique name is required for each dataset. See figure 2.6.

2.4.2 Linking Datasets, Chaining

Just like jobs can be linked to each other, datasets can link to each other too. Since datasets are built on top of jobs, this is straightforward. Assume the file `file0.txt` is imported into dataset `imp-0/default`, and that there is more data like it stored in the file `file1.txt`. The second file is imported with a link to the first dataset, see figure 2.7. The `imp-1` (or `imp-1/default`) dataset reference can now be used to access all data imported from *both* files!

Linking datasets containing related content is called *chaining*, and this is particularly convenient when dealing with data that grows over time. A good example is any kind of *log* data, such as logs of transactions, user interactions, and similar. Using chaining, new data can be added and associated with already existing data just by linking, which is a very lightweight constant time operation.

For further information on how to manually create the said datasets, visit Section 5.9.

2.4.3 Adding New Columns to a Dataset

In the previous section it was shown that datasets can be chained and thereby grow in number of rows. A dataset chain is created simply by linking one dataset to the other, so the overhead is minimal. Based on the same principles, it is equally simple to add new columns to existing datasets. Adding columns is a common operation and the Accelerator handles this situation efficiently using links.

The idea is very simple. Assume a “source” dataset to which one or more new columns should be added. A new dataset is created containing *only* the new column(s), and while creating it, the constructor is instructed to link all the source dataset’s columns to the new dataset such that the new dataset appears to contain all columns from both datasets. (Note that this linking is similar to but different from chaining.)

Accessing the new dataset will transparently access all the columns in both the new and the source dataset in parallel, making it indistinguishable from a single dataset. See Figure 2.8.

New columns are typically based on already existing ones. Then, values are written to the new columns in the new dataset while reading (iterating) from the existing columns in the source dataset. This will be discussed in detail in section 5.10




Figure 2.6: Importing `file0.txt`. The ellipse denotes a job, and the inscribed rectangle denotes a dataset. Unless a specific name is given, the name of the dataset will be `default`. The dataset can later be referenced by `imp-0/default`, or, since there is only one dataset in `imp-0`, by the name of the job that created it: `imp-0`.




Figure 2.7: Chaining the import of `file1.txt` to the previous import of `file0.txt`.

2.4.4 Multiple Datasets in a Job

Typically, a method creates a single dataset in the job directory, but there is no limit to how many datasets that could be created and stored in a single job directory. This flexibility opens up for a number of compact and efficient solutions to various common problems.

For example, one application for keeping multiple datasets in a job is the `csvimport_zip` method, which reads compressed zip archives and runs `csvimport` on each of the included files. Each file is turned into a dataset, and thus there will be one dataset in the job per file in the archive.

@@marker

2.4.5 Parallel Dataset Access and Hash Partitioning

As shown earlier in this chapter, data in datasets is stored in multiple files. Columns are stored independently, and each column is *sliced* into *slices* independent files. In this fashion, columns can be read independently without overhead, and multiple CPUs can operate on different files/slices from the same column(s) at the same time. @@@@ Datasets can be partitioned, or sliced, in different ways. One obvious way is to use round robin, where each consecutive input data row is written to the next slice, modulo the total number of slices. This leads to “well balanced” datasets with approximately equal number of rows per slice.

@@@@ Another interesting slicing alternative is to slice based on a *hash function* applied to a particular column’s values. Using this method, all rows with the same hash value in the hashed column end up in the same slice. @@@@

2.4.6 Dataset Column Types

There are large a number of useful types available for dataset columns. They include *floating* and *integer point numbers*, *Booleans*, *timestamps*, several *string types* (handling all kinds of string encodings), and *json* as well as *pickle* types for storing arbitrary data collections. Most of these types come with advanced parsers, making importing data from text files straightforward with deterministic handling of errors, overflows, and so on.




Figure 2.8: Adding one new column to the source dataset.

2.4.7 Dataset Attributes

The dataset has a number of attributes associated with it, such as shape, number of rows, column names and types, and more. An attribute is accessed like this

```
datasets = ('source',)
def synthesis():
    print(datasets.source.shape)
    print(datasets.source.columns)
```

and so on.

2.5 Iterators: Working with Data

Data in a dataset is typically accessed using an *iterator* that reads and streams one dataset slice at a time to a CPU core. The parallel processing capabilities of the Accelerator makes it possible to dispatch a set of parallel iterators, one for each slice, in order to have efficient parallel processing of the whole dataset.

This section shows how iterators are used for reading data, how to take advantage of slicing to have parallel processing, and how to efficiently create new datasets.

2.5.1 Iterator Basics

Basic iterator functionality is introduced here using an example. Assume a dataset that has a column containing movie titles named `movie`, and the problem is to extract the ten most frequently occurring movies. Here's a complete method solving this problem

```
from collections import Counter
datasets = ('source',)

def synthesis():
    c = Counter(datasets.source.iterate(None, 'movie'))
    top10 = c.most_common(10)
    print(top10)
    return top10
```

This will print the ten most common movie titles and their corresponding counts in the source dataset. The code will run on a single CPU core, because we use the single-process `synthesis()` function, which is called and executed only once. The `dataset.iterate()` (class-)method therefore has to read through all slices, one at a time, in a serial fashion, and this is reflected by the first argument to the iterator being `None`. The method also returns the variable `top10` so that it can be used by other methods.

2.5.2 Parallel Execution

It is easy to write parallel programs using the Accelerator. The fact that data in a dataset is sliced into disjoint sets and files makes parallel processing of data straightforward. Here's a slightly modified version of the program from the previous section that will now execute in parallel

```
def analysis(sliceno):
    return Counter(datasets.source.iterate(sliceno, 'movie'))

def synthesis(analysis_res):
    c = analysis_res.merge_auto()
    top10 = c.most_common(10)
    return top10
```

For larger datasets, this parallel version of the movie title counter will run much faster. Here, `.iterate()` is moved inside the `analysis()` function. This function is forked once for each slice, and the argument `sliceno` will contain an integer between zero and the number of

slices minus one. The returned value from the analysis functions will be available as input to the synthesis function in the `analysis_res` Python iterable. It is possible to merge the results explicitly, but the iterable comes with a rather magic method `merge_auto()`, that merges the results from all slices into one based on the data type. It can for example merge `Counters`, `sets`, and composed types like `sets` of `Counters`, and so on.

2.5.3 Iterating over Several Columns

Since each column is stored independently in a dataset, there is no overhead from reading a subset of a dataset's columns. In the previous section we've seen how to iterate over a single column using `.iterate()`. Iterating over more columns is straightforward by feeding a list of column names to `.iterate()`, like in this example

```
from collections import defaultdict
datasets = ('source',)

def analysis(sliceno):
    user2movieset = defaultdict(set)
    for user, movie in datasets.source.iterate(sliceno, ('user', 'movie')):
        user2movieset[user].add(movie)
    return user2movieset
```

This example creates a lookup dictionary from users to sets of movies. Note that in this case, we would like to have the dataset hashed on the `user` column, so that each user appears in exactly one slice. This will make later merging (if necessary) much easier.

It is also possible to iterate over all columns by specifying an empty list of columns or by using the value `None`.

```
...
def analysis(sliceno):
    for columns in datasets.source.iterate(sliceno, None):
        ...
```

Here, `columns` will be a list of values, one for each column in the dataset, in sorted column name order.

2.5.4 Iterating over Dataset Chains

The `iterate` function is used to iterate over a single dataset. There is a corresponding function, `iterate_chain`, that is used for iterating over chains of datasets. This function takes a number of arguments, such as

length, i.e. the number of datasets to iterate over. By default, it will iterate over all datasets in the chain.

callbacks, functions that can be called before and/or after each dataset in a chain. Very useful for aggregating data between datasets.

stop_id which stops iterating at a certain dataset. This dataset could be from *another* job's parameters, so we can for example iterate exactly over all new datasets not covered by a previous job.

range, which allows for iterating over a range of data.

The `range` options is based on the max/min values stored for each column in the dataset. Assuming that the chain is sorted, one can for example set

```
range={'ts_column': ('2016-01-01', '2016-01-31')}
```

in order to get rows within the specified range only. Using `range=` is quite costly, since it requires each row in the dataset chain with dates within the range to be checked against the range criterion. Therefore, there is a `sloppy` version that iterates over complete datasets in the chain that contains at least one row with a date within the range. This runs at full speed, and is useful, for example, to very quickly produce histograms or plots of subsets of a huge dataset.




Figure 2.9: Execution flow and result propagation in a method.

2.5.5 Job Execution Flow and Result Passing

Execution of code in a method is either parallel or serial depending on which function is used to encapsulate it. There are three functions in a method that are called from the Accelerator when a method is running, and they are `prepare()`, `analysis()`, and `synthesis()`. All three may exist in the same method, and at least one is required. When the method executes, they are called one after the other.

`prepare()` is executed first. The returned value is available in the variable `prepare_res`.

`analysis()` is run in parallel processes, one for each slice. It is called after completion of, and actually forked from `prepare()`. Common input parameters are `sliceno`, holding the number of the current process instance, and `prepare_res`. The return value for each process becomes available in the `analysis_res` variable.

`synthesis()` is called after the last `analysis()`-process is completed. It is typically used to aggregate parallel results created by `analysis()` and takes both `prepare_res` and `analysis_res` as optional parameters. The latter is an iterator of the results from the parallel processes.

Figure 2.9 shows the execution order from top to bottom, and the data passed between functions in coloured branches. `prepare()` is executed first, and its return value is available to both the `analysis()` and `synthesis()` functions. There are `slices` (a configurable parameter) number of parallel `analysis()` processes, and their output is available to the `synthesis()` function, which is executed last.

Return values from any of the three functions may be stored in the job's directory making them available to other jobs.




Figure 2.10: Execution flow of a method. The method takes optionally three kinds of parameters: `options`, `jobs`, and `datasets`.

2.5.6 Job Parameters

We've seen how completed jobs can be used as input to new jobs. Jobs are one of three kinds of input parameters that a job can take. Here the input parameters are summarised:

`jobs`, a set, or list, of identifiers to previously executed jobs;

`options`, a dictionary of options; and

`datasets`, a set, or list, of input *datasets*.

See Figure 2.10. Parameters are entered as global variables early in the method's source.

2.6 A Class Based Programming Model

The Accelerator is based on an class based paradigm. Access to the Accelerator's built in functions and parameters are typically done through a few *objects* that are populated by the running Accelerator.

Chapter 3

Basic Build Scripting

DRAFT

Build scripts are used to execute jobs and control the job flow on the Accelerator. This chapter describes the basics of job building. More advanced features, using the Urd server, are presented in chapter 7.

3.1 Build Scripts

Build scripts are stored in *method directories*. The default build script is named “`build.py`”, all other build scripts have to start with the string “`build_`”. Build scripts are executed using the `run` command. For example,

```
ax run example
```

will look for a file named `build_example.py` and execute it, while the command

```
ax run
```

by itself will execute the default build script. (The `run` command is described in section 9.5.)

The `run` command will load the build script and execute its `main()` function. This function takes a *mandatory* argument named `urd`, so a basic build script looks like this

```
def main(urd):  
    # do something here...
```

The `run` command inserts an object of the `Urd` class as the argument to the `main()` function. This `urd` object has a number of member functions and attributes useful for job building and tracking of old jobs. The `Urd` class is described in chapter 7 and in section C.7.

3.1.1 Building a Job: `urd.build()`

The `.build()` function is used to build a job from a method (i.e. source file). For example, the most simple build script that executes a method, here called `method1`, is

```
def main(urd):  
    urd.build('method1')
```

Common options for the `build` function are as follows

```
job = urd.build(  
    method,  
    options={}, datasets={}, jobs={},  
    name=None, caption=None, workdir=None,  
    **kw,  
)
```

In practice, only a few parameters are given at a time, and the only mandatory argument is the `method`. Any arguments to a job build can be added as simple *keyword arguments* (`key=value`), and the build process will automatically assign them by name to the corresponding `options`, `datasets`, or `jobs` parameter sets, assuming the naming is unique. If the key name is not unique, the argument has to be specified explicitly. Examples will be present in the following sections.

When the job is completed, `Urd` will record it using the name of the method as key, unless the `name=` is specified. The `name` parameter is particularly useful to tell jobs apart that are based on the same method. A common case would be the `csvimport` method, for example. It is also possible to assign a caption to a job, but this has no functional benefits.

When the job has been successfully built, the `build` function will return a reference of type `Job`. The `Job` class contains member functions and attributes that can be used to extract information, such as generated files or text written to `stdout`, from the job. The `Job` class is described in detail in chapter 4 and in section C.1.

Similarly, if the job to be built already exists in a configured `workdir`, the `build` function will immediately return a `Job` object corresponding to the existing job without executing anything.

3.1.2 Connecting Jobs

Jobs may be connected, or linked, together so that a new job may use the output generated by previous jobs. In this way, jobs could share and re-use previous results, and complicated tasks can be broken down into several smaller jobs. A link is created by feeding the job object from an existing job into the `.build()`-call of a new job, like in this example

```
def main(urd):
    job_import = urd.build('csvimport', filename='inputfile.txt')
    job_process = urd.build('process', source=job_import)
```

In the example above, the first job, `csvimport`, imports the file “`inputfile.txt`”. The second job, `process`, takes the imported dataset as input for further processing.

3.1.3 Replaying Build Scripts

When the example build script from the previous section is run, both the `csvimport` and the `process` jobs will be built. But what happens if the same build script is run a second time? Remember now that the Accelerator stores all jobs in associated `workdir(s)`. If there has been no changes to the code since the last run, the Accelerator will immediately find the job reference to the `csvimport` without needing to execute it again. The reference will be input to the `.build()` call of the second method, and since the Accelerator has seen this call before too, it will immediately look up the reference to this job as well, instead of executing it. A second run of the build script will only take a fraction of a second to execute, but it will still return all job references.

On the other hand, if something has been modified, such as a method’s source code or any of the input parameters, the affected job(s) will be re-executed. For example, assume that the `csvimport` build call is modified to import “`anotherfile.txt`” instead. This will cause the `csvimport` method to be executed again, leading to a new job with a new `jobid`. This job reference is input to the `process` job, causing it to be re-executed too. Here, both jobs were re-executed, since they both were affected by the modification, but in general, only those jobs affected by the modification will be re-executed!

A successful “replay” of a build script ensures the integrity and dependencies of all involved calculations. If there are no changes, the same result remains. If, however, some of the code has been modified, the Accelerator will compute new jobs to reflect the new situation. The result may be different, and the user is notified.

3.2 Working with Build History: `urd.joblist`

Information about previously executed jobs is stored in the `urd.joblist` variable. This variable is of type `JobList`, which is basically a standard ordered Python list with some additional features for searching, profiling and pretty-printing. The `JobList` class is further explained in section C.3.

3.2.1 Printing a JobList: `urd.joblist.pretty`

Create a `JobList` and pretty-print it

```
def main(urd):
    job1 = urd.build('first')
    job2 = urd.build('second', first=job1)
    print(urd.joblist.pretty)
```

which results in

```
JobList(
  [ 0] first : TEST-38
  [ 1] second : TEST-39
)
```

(The actual `jobids` will most likely be different.) The name in the `joblist` is either the name of the method, or, if present, the name given explicitly using the `urd.build(name=)` option.

3.2.2 Finding Jobs in a Joblist

There are several ways to extract jobs or list of jobs from a `joblist`.

Using `find` to Extract Jobs to a New Joblist

The `find()` function finds matching jobs and returns them in a new `JobList`. For example,

```
j1 = urd.joblist.find('csvimport')
```

will create a `JobList` of all `csvimport` jobs in `urd.joblist`.

Using `get` to Potentially Extract a Single Job

The `get` function will return a job reference to the most recent matching job. For example,

```
job = urd.joblist.get('csvimport')
```

If no matching job is found, `get` will return `None`.

Using Square Brackets to Extract a Single Job

Accessing jobs directly with a key like this

```
job = urd.joblist['csvimport']
```

is similar to `get`, but will return an error if a matching job is not found.

3.2.3 Indexing and Slicing a JobList

Since the `JobList` class is an extension of the `(list)` class, it is possible to do some `list`-like operations on `joblist` objects.

A common special case is to provide an integer to the `.get()` function. This will cause a lookup of the item with the corresponding index in the list. It is often used like this

```
latest_job = urd.joblist.get(-1)
```

as a way to return the last item in the `joblist`, should it exist, or `None`.

It is also possible to index and slice like in the following examples

```
joblst = j1[3]
```

or

```
joblst = j1[-2:]
```

3.3 Configuration Information: `urd.info`

The dictionary `urd.info` contains configuration information from the Accelerator server. In particular, it contains at least these fields

name	description
<code>slices</code>	Configured number of slices.
<code>urd</code>	An URL to the Urd server.
<code>result_directory</code>	see section B.4.
<code>input_directory</code>	see section B.4.

3.4 Summary

In a build script, the `urd` object has functionality for building and retrieving jobs. A job is built using `urd.build()`, and references to all built jobs are stored in `urd.joblist`. These references could be fed as input parameters to new jobs so that the output from one job could be used as input by another. The `urd.joblist` variable is basically of type `list`, but with extra functionality to conveniently find already existing jobs.

DRAFT

Chapter 4

Jobs

DRAFT

4.1 Definitions

4.1.1 Methods and Jobs

In general, doing a computation on a computer follows the following equation

$$\text{source code} + \text{input data and parameters} + \text{execution time} \rightarrow \text{result}$$

In the Accelerator context, the notation is as follows

$$\text{method} + \text{input data} + \text{input parameters} + \text{execution time} \rightarrow \text{job}$$

where the **method** is the source code, and the **job** is a directory containing

- any number of output files created by the running method, as well as
- a number of job meta information files containing all information that was needed to run the job in the first place.

The exact contents of the job directory will be discussed in section 4.1.3, but note here that the directory contains everything needed to run a specific piece of code with inputs and source code, *as well as* any output generated during execution.

Computing a job is denoted job *building*. Jobs are **built** from methods. When a job has been built, it is *static*, and cannot be altered or removed by the Accelerator. Jobs are built either by

- a *build script*, see chapter 7, or
- by a method, using *subjobs*, see section 4.10

The following figure illustrates how a job `example-1` is built from the method `a_method.py`. The job is stored in the `example` work directory. The job identifier (in this case `example-1`) is always unique so that it can be used as a reference to that particular job.




Figure 4.1: When a method is built, a job directory is created in the target work directory, containing files with all data and meta information regarding the job.

4.1.2 Jobids

A **jobid** is a string that can be used as a reference to a job. This unique string is created by appending an incrementing integer to the name of the work directory in which the job is stored. In the example above, the job is uniquely identified by the string `example-0`.

4.1.3 Work Directories and Job Directories

A successful build of a method results in a new job directory on disk. The job directory will be stored in the current *workdir* (i.e. work directory). *This section is for reference only, since jobs and datasets are much more conveniently inspected using the `ax job` and `ax ds` shell commands.*

The structure of a job directory is similar to the following, assuming the current *workdir* is `test`, and the current *jobid* is `test-0`.

```
workdirs/test/
  test-0/
    setup.json
    method.tar.gz
    result.pickle
    post.json
    OUTPUT/
    default/
```

The following table shows examples of files commonly found in a job directory.

name	description
setup.json	Contains information about the job build, including name of method, input parameters, and, after execution, some profiling information.
post.json	Contains profiling information, and is written only if the job builds successfully.
method.tar.gz	All source files, i.e. the method's source and any <code>depend_extras</code> are stored in this gzipped tar-archive.
result.pickle	The return value from <code>synthesis()</code> stored in the Python "pickle" format.
default/	If the job contains datasets, these will be stored in directories, such as for example <code>default/</code> , in the root of the job directory.
datasets.txt	List of all datasets in job in a human readable format.
OUTPUT/	Any output to <code>stdout</code> and <code>stderr</code> will be stored in the <code>OUTPUT/</code> directory.

4.1.4 The Job and CurrentJob Convenience Wrappers

In order to simplify access to job directory data, common job data operations are made available by the `Job` class. There is also an extended version of this, called the `CurrentJob` class, that also contains information and helper functions to a *running* job. See section C.1 for details about these classes.

4.2 Python Packages

Methods are stored in standard Python packages, i.e. in directories that are

- reachable by the Python interpreter, and
- contain the (perhaps empty) file `__init__.py`.

Packages are by default located in the same directory as the Accelerator's configuration file, but can be changed by setting the `package_directory` configuration option in the Accelerator's configuration file.

There are two ways, to define a package, either using `auto-discover`, or using `methods.conf` files, where the first is recommended for new users and the latter strongly recommended for production systems and projects with more strict requirements.

4.2.1 A New Package using auto-discover

In `auto-discover`, methods are automatically detected by the Accelerator, and anything detected can be executed. Create a new directory and make it a Python package like this

```
% mkdir <dirname>
% touch <dirname>/__init__.py
```

Then, add the string

```
<dirname> auto-discover
```

to the `method packages` section of the configuration file, see section B.4

4.2.2 A New Package using `methods.conf`

Create a new package directory

```
% mkdir <dirname>
% touch <dirname>/__init__.py
% touch <dirname>/methods.conf
```

and add the string

```
<dirname>
```

to the `method packages` section of the configuration file, see section B.4. The first two lines create a Python package, and the third line adds the file `methods.conf`, which is required in “strict” mode.

For security reasons, the Accelerator now only looks for packages explicitly specified in the configuration file using the “`method packages`” assignment. Again, see chapter B.4 for detailed information about the configuration file.

4.3 Method Source Files

Method source files are stored in Python packages as described in the previous section. The Accelerator searches all reachable packages for methods to execute, and therefore *method names need to be globally unique!* In order to reduce risk of executing the wrong file, these limitations apply to methods:

1. the method name must be *globally* unique, i.e. there can not be a method with the same name in any other method directory visible to the Accelerator;
2. for a method file to be accepted by the Accelerator, the filename has to start with the prefix “`a_`”; and
3. *unless* `auto-discover` is enabled (see previous section), the method name, *without* the “`a_`” prefix, must be present on a separate line in the `methods.conf` file for the package, see section 4.3.2.

4.3.1 Creating a New Method

In `auto-discover` mode, it is straightforward. Any file named `a_name.py` will correspond to a method named `name`. Here’s how to add a method:

1. Create the method in a package visible to the Accelerator using an editor. Make sure the filename is `a_name.py` if the method’s name is `name`.
2. Add the method name `name` (without the prefix “`a_`” and suffix “`.py`”) to the `methods.conf` file in the same method directory as where the source file is stored. See section 4.3.2. This is *not* needed if `auto-discover` is enabled, but it can still be used to specify a particular Python interpreter for the method, see next section.
3. (Remember to make sure that the method directory is defined in the Accelerator’s configuration file.)

4.3.2 Making Methods Executable: `methods.conf`

The file `methods.conf` provides an easy way to specify and limit which methods (source files) that can be executed, and optionally, which Python interpreter that should be used for each of them. Methods not specified in `methods.conf` cannot be executed, unless `auto_discover` is enabled. (Execution control is a must in production environments to prevent accidental execution of “forgotten” test files etc.)

The `methods.conf` is a plain text file with one entry per line. Any characters from a hash sign (“#”) to the end of the line is considered to be a comment. It is permitted to have any number of empty lines in the file. Available methods are entered first on a line by stating the name of the method, without the `a_` prefix and `.py` suffix.

The method name can optionally be followed by one or more whitespaces and a name specifying the actual Python interpreter that will be used to execute the method. This is particularly useful to have individual methods run from individual virtual environments. Thus, each method can run from its own virtual environment, with its own dependencies. A list of valid Python interpreters is defined in the configuration file using the key “`interpreters`”, see section B.4.

The default interpreter is selected if the second field is left empty, where default corresponds to the one that the currently running Accelerator server is using. The Accelerator and its `standard_methods` library are compatible with both Python 2 and Python 3.

Here is an example `methods.conf`

```
# this is a comment
test2                # will use default Python
test3                py3 # py3 as specified in accelerator.conf
testx                tf  # a Tensorflow virtual env defined in accelerator.conf

```

This file declares three methods corresponding to the files `a_test2.py`, `a_test3.py`, and `a_testx.py`. These are the only methods that can be built in this method package. Note that some methods are using virtual environments specified in the Accelerator’s configuration file.

4.4 Job Building or Job Recycling

Since the Accelerator keeps track of a job’s dependencies and results, it can in an instant determine if a job to be built has been built before. If the job has been built before, the Accelerator will immediately return a job instance to the existing job. Otherwise, the job will first be built, and then a job instance will be returned.

4.4.1 Job Already Built Check

As shown in chapter 3, jobs are built using the `.build()` function, like this

```
def main(urd):
    urd.build('themethod', parameter1=..., ...)
```

Prior to building a method, the Accelerator checks if an equivalent job has been built in the past. This check is based on two things:

1. the output of a hash function applied to the method source code, and
2. the method’s input parameters.

The hash value is combined with the input parameters and compared to all jobs already built. Only if the hash and input parameter combination is unique, i.e. it has not been seen before, will the method be executed. The `.build()`-function returns an instance of type `Job`. To the caller, it is not apparent if the job was just built or if it was built at an earlier time.

4.4.2 Depend on More Files: `depend_extra`

A method may import code located in other files, and such files can be included in the build check hash calculation as well. This will ensure that a change to an imported file will indeed force a re-execution of the method if a build is requested. Additional files are specified in the method using the `depend_extra` list, as for example:

```
from . import my_python_module

depend_extra = (my_python_module, 'mystuff.data',)
```

As seen in the example, it is possible to specify either Python module objects or filenames relative to the method's location.

If the Accelerator suspects that a `depend_extra`-statement is missing, it will suggest adding it by printing a message in the output log like this:

```
=====
WARNING: dev.a_test should probably depend_extra on myfuncs
=====
```

The point of this is to make the user aware that the method depends on additional files that are currently not taken into account in the build check hashing. The warning is removed by putting the `myfuncs` file in a `depend_extra` list of the `test` method.

4.4.3 Avoiding Rebuild: `equivalent_hashes`

A change to a method's source code will cause a new job to be built upon running `.build()`, but sometimes it is desirable to modify the source code *without* causing a re-build. This happens, for example, when new functionality is added to an existing method, and re-computing all jobs is not an option. If functionality remains the same, existing jobs strictly do not need to be re-built. For this situation, there is an `equivalent_hashes` dictionary that can be used to *manually* specify which versions of the source code that are equivalent. The Accelerator helps creating this dictionary, if needed. This is how it works.

1. Find the hash `<old_hash>` of the existing job in that job's `setup.json`.
2. Add the following line to the updated method's source code

```
equivalent_hashes = {'whatever': (<old_hash>,)}
```

3. Run the build script. The server will print something like

```
=====
WARNING: test_methods.a_test_rechain has equivalent_hashes,
but missing verifier <current_hash>
=====
```

4. Copy `<current_hash>` into the `equivalent_hashes`:

```
equivalent_hashes = {<current_hash>: (<old_hash>,)}
```

This line now tells that `current_hash` is equivalent to `old_hash`, so if a job with the old hash exists, the method will not be built again. Note that the right part of the assignment is actually a list, so there could be any number of equivalent versions of the source code. From time to time, this has been used during development of the Accelerator's `standard_methods`, when new features have been added that do not interfere with existing use.

4.4.4 Forcing a Rebuild

It is possible to force build a job, i.e. build a job that already exists. This is mainly implemented for test code purposes, and should be avoided in most development methodologies. This is how to do a force build

```
urd.build(method, force_build=True)
```

Without forcing a re-build, all jobs in a work directory are unique. But it is still possible to have duplicate jobs by adding/removing workdirs and building the same job to different workdirs. There are no inherent problems with multiple job instances, but there is typically no need for it, unless in testing. Another way to force a job to be re-built is by supplying a timestamp as input option.

4.5 Invoking the Python Debugger

Although the Accelerator is a client-server application, it is possible to invoke the Python Debugger at any time. The only thing needed is to start the server with a flag like this

```
ax server --debuggable
```

When this flag is set, methods will be executed in the same session as the server with that terminal as `stdin`. This means that it is possible to write for example

```
breakpoint()
```

anywhere in a method and communicate with the debugger in the server terminal window. Note that using this flag will kill the server if the method fails. See section 9.4

4.6 Method Execution

Methods are executed using the `build()` call, either from a `build script`, or from another method as a `subjob`. Methods typically takes input parameters, and they may generate return values and produce output files as well as output to `stdout` and `stderr`.

4.6.1 Input Parameters

There are three kinds of method input parameters assigned by the `build()` call: `jobs`, `datasets`, and `options`. These parameters are stated early in the method source code and are *global*, meaning that they do not need to be included as parameters to the functions in a method. Here is an example parameter set

```
jobs = ('accumulated_costs',)
datasets = ('transaction_log', 'access_log',)
options = dict(length=4)
```

The input parameters are populated by the builder when the `run` command is executed. Section 4.9 and 4.11 provide detailed descriptions of all parameters.

4.6.2 Functions Reserved for Execution Flow

During execution, methods are not run from top to bottom. Instead, there are three reserved functions that are called by the method dispatcher controlling the execution flow. These functions are

```
prepare(),
analysis(), and
synthesis().
```

4.6.3 Execution Order

The three functions `prepare()`, `analysis()`, and `synthesis()` are called one at a time in that order. `prepare()` and `synthesis()` execute as single processes, while `analysis()` provides parallel execution. None of them is mandatory, but at least one must be present for the method to execute. It is discouraged to use `prepare()` only. Use `synthesis()` for single-process execution.

4.6.4 Function Arguments

There are some optional arguments that can be passed into the executing functions `prepare()`, `analysis()`, and `synthesis()` at run time. All functions have access to these

- `job`, which is an instance of the current job, and
- `slices`, an integer holding the total number of slices.

while only `analysis()` has access to

- `sliceno`, which provides a unique number to each parallel `analysis()`-process. This is also a **mandatory** argument to `analysis()`.

The `job` instance contains information and helper functions regarding the current job. The object is of type `CurrentJob`, which is an extension of the `Job` class used for job instances that are not in the execution stage.

The `analysis()` function (and only the `analysis()` function) takes the mandatory argument `sliceno`, which is an integer between zero and the total number of slices minus one. This is the unique identifier for each `analysis()` process, and it is commonly used when accessing sliced datasets or doing other parallel processing tasks, see for example chapter 6 for its use in dataset iterators.

4.6.5 Parallel Processing: The `analysis()` Function and Slices

The number of parallel analysis processes is set by the `slices` parameter in the Accelerator's configuration file. The input parameter `sliceno` to the `analysis()` function is the unique identifier for each parallel function call, and its value is in the range from zero to the number of slices minus one. When processing Accelerator *datasets*, the idea is that each dataset slice should have exactly one corresponding `analysis()` process, so that all the slices in the dataset can be processed in parallel.

4.6.6 Return Values

Return values may be passed from one function to any function that will execute later. To be specific, what is returned from `prepare` is called `prepare_res`, and can be used as input argument to `analysis()` and `synthesis()`. Furthermore, the return values from `analysis()` are available as `analysis_res` in `synthesis()`. The `analysis_res` variable is an iterator, yielding the results from each slice in turn. Finally, the return value from `synthesis()` is stored permanently in the job directory using the name “`result.pickle`”. Note that `prepare_res` and `analysis_res` are set to `None` if nothing is returned. Here is an example of return value passing

```
options = dict(length=4)

def prepare():
    # options is a global variable
    return options.length * 2

def analysis(sliceno, prepare_res):
    return prepare_res + sliceno

def synthesis(analysis_res, prepare_res):
    return sum(analysis_res) + prepare_res
```

Note that when a job completes, it is not possible to retrieve the results from `prepare()` or `analysis()` anymore. Only results from `synthesis()` are kept. Creating persistent intermediate files is the topic of section C.1.17, however.

4.6.7 Automatically Merging Results from `analysis()`

It is common that results from different slices needs to be merged together. The Accelerator provides a relatively general function for merging sliced data structures. Consider this example

```

# create a set of all users
datasets = ('source',)

def analysis(sliceno):
    return(set(datasets.source.iterate(sliceno, 'user')))

def prepare(analysis_res):
    return analysis_res.merge_auto()

```

Here, each `analysis()` process creates a set of users seen in that *slice* of the `source` dataset. In order to create a set of all users in the dataset, all slice-sets have to be merged. Merging can be implemented using for example a `for`-loop, but the actual merging operation is dependent of the actual data type, and writing merging functions is error prone. (So we've been told!) Therefore, `analysis_res` has a function called `merge_auto()`, that is recommended for merging. This function can merge most data types, and even merge container variables in a recursive fashion. For example, a variable defined like this

```
h = defaultdict(lambda: defaultdict(set))
```

(a dict of dicts of sets) is straightforward to merge using `merge_auto()`. The function works on many data types and is less error-prone than writing special mergers every time they are needed.

4.6.8 Standard Out and Standard Error

Anything sent to `stdout` or `stderr` during job execution will be sent *both* to the terminal in which the Accelerator server was started, *and* to a file in the current job directory. This covers, for example, anything output from Python's `print()`-function. Thus, the job directory contains a complete log of all printed output from the execution phase!

Output is collected in the job directory in a subdirectory named `OUTPUT`, and it is made available using the `.output()` function, see section 4.7.7. The `OUTPUT` directory is created *only* if anything was output from the job to `stdout` or `stderr`, otherwise it does not exist. Inside the directory there may be files like this

```

job-x/
  OUTPUT/
    prepare      # created if output in prepare()
    synthesis    #                synthesis()
    0            #                analysis() slice 0
    3            #                3

```

No empty files will be created.

4.7 The Job and CurrentJob Classes

The `Job` and `CurrentJob` classes provide functionality for easy access to data and datasets stored in a job directory. (Datasets will be covered in chapter 5). The `CurrentJob` is an extension of `Job` that adds special functions that are useful to a method during execution. This section provides a taste of the most common operations that are provided. See section C.1 for a complete list of the functionality.

Instances of these two classes are used extensively in Accelerator projects. In a build script every reference to a job, such as the return value of the `.build()` function or any job retrieval using the Urd database are of type `Job`. Any job passed as input parameter to a `.build()`-call will appear as a `Job` instance inside the running method. Instances of the `CurrentJob` class are provided when asking for a `job` input parameter in `prepare()`, `analysis()`, or `synthesis()`.

4.7.1 Writing and Reading Serialised Data

Data structures may be serialised and written to disk using `job.save()` and `job.json_save()`, with corresponding `.load()` and `.json_load()` functions, where the first writes a Python

“pickle” file, and the latter uses JSON encoding. Here is an example of how to write files in a running job

```
def synthesis(job):
    job.save('a string to be written', 'stringfile')
    job.json_save(dict(key='value'), 'jsonfile')
```

The corresponding `job.load()` and `job.json_load()` functions can be called *both* in methods *and* build scripts. For example

```
jobs = ('anotherjob',)

def synthesis():
    jobs.anotherjob.load('stringfile')
```

will load a file from another job into the currently running method, while

```
def main(urd):
    job = urd.build('example')
    x = job.load('thefile')
```

will load data stored by the `example` method using the filename `thefile.pickle` into the build script.

4.7.2 Writing and Reading Serialised Data in Parallel

If data is read and written in the parallel `analysis()`-function, the argument `sliceno=` may be used to write one file for each slice. For example

```
def analysis(sliceno, job)
    data = ...
    job.save(data, 'filename', sliceno=sliceno)
```

Similarly, another job can then read one of these files per slice as follows

```
def analysis(sliceno):
    data = jobs.anotherjob.load('filename', sliceno=sliceno)
```

Writing “sliced” data results in n files on disk, where n is equal to the number of slices set in the configuration file. Each filename is extended with a human readable number that corresponds to the slice that the file’s data belongs to.

4.7.3 General File Access

The `.open()` function corresponds to the built in `open()` with the addition that it cannot write to completed jobs, and that files written using it are book-kept in the job. It has to be used as a context manager, i.e. using the `with` statement, for example like this

```
def synthesis(job):
    with job.open('filename', 'wb') as fh:
        fh.write(...)
```

4.7.4 Accessing a Job’s Return Value

The default behaviour of a job instance’s `.load()` function is to read the return value from the job’s `synthesis()` function, like this

```
def main(urd):
    job = urd.build('example')
    x = job.load()
```

This works both in build scripts and inside methods, and is a convenient way to access data generated by a job.

4.7.5 Accessing a Job's Datasets

Using the Job class, it is straightforward to access datasets in other jobs. For example

```
def main(urd):
    job = urd.build(...)

    # This will print a list of all dataset instances in the job.
    print(job.datasets())

    # This will return a dataset instance of the job/training dataset.
    ds = job.dataset('training')
```

This works both in running methods and in build scripts

4.7.6 Accessing a Job's Options and Parameters

There are two sources of parameters to a running method,

parameters from the caller, i.e. the `.build()`-call, and

parameters assigned by the Accelerator when the job starts building.

All these parameters are available using `job.params`. For example

```
jobs = ('anotherjob',)
```

```
def synthesis():
    print(jobs.anotherjob.params.options)
```

will print the options dictionary that was fed to the `anotherjob` at build time, for example

```
{'message': 'Hello world!'}
```

A complete print of a job's `.params` may look like this

```
{
  "starttime": 1602061144.081299,
  "endtime": 1602061147.2101562,
  "exectime": {
    "analysis": 0,
    "per_slice": [],
    "prepare": 0,
    "synthesis": 3.111,
    "total": 3.111
  },
  "caption": "",
  "hash": "9189b775e190826f3dc6ea85ea252a9e3d647185",
  "jobid": "beast-325",
  "method": "plot_walk_narrowbeams",
  "package": "dev",
  "seed": 3621427863964846452,
  "slices": 4,
  "version": 3,
  "versions": {
    "accelerator": "2020.10.3.dev1",
    "python": "3.5.2 (default, Jul 17 2020, 14:04:10) \n[GCC 5.4.0 20160609]",
    "python_path": "/home/eaenbrd/checkout/project_beast.acc/venv/bin/python3"
  },
  "options": {
    "gnbdir": 0.3839724354387525,
    "gnbpos": [
```

```

        57.71525125357654,
        12.83168
    ]
},
"datasets": {
    "source": "beast-222"
},
"jobs": {
    "background": "beast-4"
}
}

```

and a description of its keys

name	description
package	Python package for this method
method	name of this method
jobid	jobid of this job
starttime	start time in epoch format
endtime	end time in epoch format
exectime	various exec times
caption	a caption
slices	number of slices of current Accelerator configuration
seed	a random seed available for use ¹
hash	source code hash value
versions	Accelerator and Python versions
options	input parameter
datasets	input parameter
jobs	input parameter

¹ The Accelerator team recommends *not* using `seed`, unless non-determinism is actually a goal.

Note how well covered input parameters and settings are, all the way down to the specific Python interpreter.

4.7.7 Accessing Job Output

Anything written to `stderr` or `stdout` during job execution is available using the `.output()` function. Here is an example

```

def main(urd):
    job = urd.build('example')
    print(job.output())

```

With no argument, the `.output()` function returns all output. Particular parts of the output can be selected using the options, `'prepare'`, `'analysis'`, `'synthesis'`, or a digit specifying a particular slice.

4.7.8 Reading Post Data

The `.post` attribute contains information such as `starttime`, execution time (per function and slice), written files and subjobs for a job. For example

```
def main(urd):
    job = job.build('example')
    print(job.post.exectime)
```

4.8 From Jobs to Datasets and Back

Sometimes it is necessary to find the job that created a particular dataset, or access one of the other datasets in a job given a certain dataset.

4.8.1 From Dataset to Job

```
job = ds.job
```

4.8.2 From Job to Dataset

```
ds = job.dataset("datasetname")
```

4.8.3 From Dataset to Dataset (in same Job)

```
ds = ds.job.dataset("datasetname")
```

4.9 Method Input Parameters

There are three kinds of method input parameters assign by the `build` call: `jobs`, `datasets`, and `options`. These parameters are stated early in the method source code, such as for example

```
jobs = ('accumulated_costs',)
datasets = ('transaction_log', 'access_log',)
options = dict(length=4)
```

The input parameters are populated by the builder using name matching when the job build is initiated, see chapter 7 for more information.

The `jobs` parameter list is used to input references to other jobs, while the `datasets` parameter list is used to input references to datasets. These parameters are populated by the build call.

The `options` dictionary, on the other hand, is used to input any other type of parameters to be used by the method at run time. Options does not necessarily have to be populated by the build call. Instead, “default values” may be used as “global constants” in the method. An option assigned by the build call will override the default assignment.

Note that `jobs` and `datasets` are tuples (or lists or sets), and a single entry has to be followed by a comma as in the example above, while `options` is a dictionary. Individual elements of the input parameters may be accessed inside the method using dot notation like this

```
jobs.accumulated_cost
datasets.transaction_log
options.length
```

Each of these parameters will be described in more detail in following sections.

4.9.1 Input Jobs

The `jobs` parameter is a tuple of job references linking other jobs to this job. In a running job, each item in the `jobs` tuple is of type `Job`, and it is possible to used them directly as references to corresponding jobs. All items in the `jobs` tuple must be assigned by the builder to avoid run time errors.

If the number of input jobs is not constant or known beforehand, they can be represented as a list, like in this example

```
jobs = ('source', ['alistofjobs'],)
```

where `source` is a single job reference, whereas `alistofjobs` is a list of job references.

4.9.2 Input Datasets

The `datasets` parameter is a tuple of links to datasets. In a running job, each item in the `datasets` variable is of type `Dataset`. The `Dataset` class is further described in chapter 5. All items in the `datasets` tuple must be assigned by the builder to avoid run time errors.

Similar to `jobs`, one can represent an unknown number datasets using a list, like this

```
datasets = ('source', ['alistofdatssets'],)
```

where `source` is a single dataset, whereas `alistofdatssets` is a list of datasets.

4.9.3 Input Options

The `options` parameter is of type `dict` and is used to pass various information from the builder to a job. This information could be integers, strings, enumerations, sets, lists, and dictionaries in a recursive fashion, with or without default values. Assigning defined options from a build call is not necessary, and an assignment will override the “default” specified, if any. Options are specified like in this example

```
options = dict(key=value, ... ) # or
options = {key: value, ...}
```

Options are straightforward to use and quite flexible. A formal overview is presented in section 4.11.

4.9.4 A Specific File From Another Job: JobWithFile

Any specific file from an existing job can be input to a new job at build time using `job.withfile()`. Here is an example

```
def main(urd):
    job = urd.build('example4')
    urd.build(
        'example5',
        firstfile=job.withfile('myfile1', sliced=True),
        secondfile=job.withfile('myfile2'),
    )
```

Inside the method, the option part is defined like this

```
from accelerator import JobWithFile
options=dict(firstfile=JobWithFile, secondfile=JobWithFile)
```

The `.withfile()` function requires a filename and takes two optional arguments: `sliced` and `extras`. The `extras` argument is used to pass any kind of information that is helpful when using the specified file, and `sliced` tells that the file is stored in parallel slices. (Creating sliced files is described in section 4.7.2.)

Here’s how to use the `JobWithFile` object in a running job. First, loading a file is done using `.load()`, like this

```
from accelerator import JobWithFile
options=dict(firstfile=JobWithFile, secondfile=JobWithFile)

def analysis(sliceno):
    print(options.firstfile.load(sliceno=sliceno))

def synthesis():
    print(options.secondfile.load())
```

The `.load()` function assumes that the file is in Python Pickle format. There is also an `.json_load()` function for JSON-files. To get the full filename to the file, use `.filename()`

```
print(options.firstfile.filename(sliceno=3))
print(options.secondfile.filename())
```

And finally, there is also a wrapper around `open()`, so it is possible to do

```
with(options.firstfile.open(), 'rb') as fh:
    data = fh.read()
```

4.10 Subjobs

Jobs may launch jobs, i.e. a running job may build other jobs in a recursive manner. As always, if the jobs have been built already, they will be linked in immediately. If the build of a subjob fails, the building job will be invalidated. Subjobs are built like in this example

```
from accelerator import subjobs

def prepare():
    subjobs.build('count_items', options=dict(length=3))
```

It is possible to build subjobs in `prepare()` and `synthesis()`, but *not* in `analysis()`. The `subjobs.build()` call uses the same syntax as `urd.build()` described in chapter 7, so input parameters like `options`, `datasets`, `jobs`, and `caption` are available for subjobs too. Similarly, the return value from a subjob `build()` is a job instance corresponding to the built job.

There are three catches, though.

1. Dataset instances to datasets created in subjobs will not be explicitly available to the “top level” build script. The workaround is to link the dataset to the building method like this

```
from accelerator import subjobs
def synthesis():
    job = subjobs.build('create_a_dataset')
    ds = job.dataset(<name>)
    ds.link_to_here(name=<anothername>)
```

with the effect that the building job will act like a dataset, even though the dataset is actually created and stored in the subjob. The `name` argument is optional, the `name default` is used if left empty, since this is the default dataset name.

It is also possible to override the dataset’s *previous dataset* using the `override_previous` option, which takes a dataset reference (or `None`) to be the new `previous`.

```
ds.link_to_here(name='thename', override_previous=xxx)
```

The `link_to_here` call returns the new dataset instance.

2. Subjobs are not visible in build scripts, and do not show up in for example `urd.joblist`.

There is a limit to the recursion depth of subjobs, to avoid creating unlimited number of jobs by accident. The limit can be tweaked by modifying the source code, if necessary.

4.11 Formal Option Rules

This section covers the formal rules for the `options` parameter.

1. Typing may be specified using the class name (e.g. `int`), or as a value of the desired type (e.g. the number `3`). See this example

```
options = dict(
    a=3,      # typed to int
    b=int,    #          int
    c=3.14,   #          float
    d='',     #          str
)
```

Values will be default values, and this is described thoroughly in the other rules.

2. An input option value is required to be of the correct type. This is, if a type is specified for an option, this must be respected by the builder. Regardless of type, `None` is always accepted.
3. An input may be left unassigned, unless
 - the option is typed to `RequiredOption()`, or
 - the option is typed to `OptionEnum()` without a default.

So, except for the two cases above, it is not necessary to supply option values to a method at build time.

4. If typing is specified as a value, this is the default value if left unspecified.
5. If typing is specified as a class name, default is `None`.
6. `None` is always a valid input unless
 - `RequiredOption()` without `none_ok=True`
 - `OptionEnum()` without `none_ok=True`

This means that for example something typed to `int` can be overridden by the builder by assigning it to `None`. Also, `None` is also accepted in typed containers, so a type defined as `[int]` will accept the input `[1, 2, None]`.

7. All containers can be specified as empty, for example `{}` which expects a `dict`.
8. Complex types (like `dicts`, `dicts of lists of dicts`, ...) never enforce specific keys, only types. For example, `{'a': 'b'}` defines a dictionary from strings to strings, and for example `{'foo': 'bar'}` is a valid assignment.
9. Containers with a type in the values default to `None`. Otherwise the specified values are the default contents. Example

```
options = dict(
    x=dict,      # will be None as default, accepts any dict
    y={'a': 'b'}, # will be {'a': 'b'} as default, accepts dict of str
    z={str: str}, # will be {} as default, accepts dict of str
)
```

The following sections will describe typing in more detail.

4.11.1 Options with no Type

An option with no typing may be specified by assigning `None`.

```
options = dict(length=None) # accepts anything, default is None
```

Here, `length` could be set to anything.

4.11.2 Scalar Options

Scalars are either explicitly typed, as

```
options = dict(length=int)    # Requires an int value or None
```

or implicitly with default value like

```
options = dict(length=3)     # Requires an int value or None,  
                             # default is 3 if left unassigned
```

4.11.3 String Options

A (possibly empty) string with default value `None` is typed as

```
options = dict(name=str)    # requires string or None, defaults to None
```

A default value may be specified as follows

```
options = dict(name='foo')  # requires string or None, provides default value
```

And a string required to be specified and non-empty as

```
from accelerator import OptionString  
options = dict(name=OptionString)    # requires non-empty string
```

In some situations, an example string is convenient

```
from accelerator import OptionString  
options = dict(name=OptionString('bar')) # Requires non-empty string,  
                                           # provides example (NOT default value)
```

Note that “bar” is not default, it just gives the programmer a way to express what is expected.

4.11.4 Enumerated Options

Enumerations are convenient in a number of situations. An option with three enumerations is typed as

```
# Requires one of the strings 'a', 'b' or 'c'  
from accelerator import OptionEnum  
options = dict(foo=OptionEnum('a b c'))
```

and there is a flag to have it accept `None` too

```
# Requires one of the strings 'a', 'b', or 'c'; or None  
from accelerator import OptionEnum  
options = dict(foo=OptionEnum('a b c', none_ok=True))
```

A default value may be specified like this

```
# Requires one of the strings 'a', 'b' or 'c', defaults to 'b'  
from accelerator import OptionEnum  
options = dict(foo=OptionEnum('a b c').b)
```

(The `none_ok` flag may be combined with a default value.) Furthermore, the asterisk-wildcard could be used to accept a wide range of strings

```
# Requires one of the strings 'a', 'b', or any string starting with 'c'  
options = dict(foo=OptionEnum('a b c*'))
```

The example above allows the strings “a”, “b”, and all strings starting with the character “c”.

4.11.5 List and Set Options

Lists are specified like this

```
# Requires list of intable or None, defaults to empty list
options=dict(foo=[int])
```

Empty lists are accepted, as well as `None`. In addition, `None` is also valid inside the list. Sets are defined similarly

```
# Requires set of intable or None, defaults to empty set
options=dict(foo={int})
```

Here too, both `None` or the empty `set` is accepted, and `None` is a valid set member.

4.11.6 Date and Time Options

The following date and time related types are supported:

```
datetime,
date,
time, and
timedelta.
```

A typical use case is as follows

```
# a datetime object if input, or None
from datetime import datetime
options = dict(ts=datetime)
```

and with a default assignment

```
# a datetime object if input, defaults to a datetime(2014, 1, 1) object
from datetime import datetime
options = dict(ts=datetime(2014, 1, 1))
```

4.11.7 More Complex Stuff: Types Containing Types

It is possible to have more complex types, such as dictionaries of dictionaries and so on, for example

```
# Requires dict of string to string
options = dict(foo={str: str})
```

or another example

```
# Requires dict of string to dict of string to int
options = dict(foo={str: {str: int}})
```

As always, containers with a type in the values default to empty containers. Otherwise, the specified values are the default contents.

4.12 Jobs - a Summary

The concepts relating to Accelerator jobs are fundamental, and this section provides a shorter summary about the basic concepts.

1. Data and metadata relating to a job is stored in a job directory.
2. Job objects are wrappers around job directories, providing helper functions.

The files stored in the job directory at dispatch are complete in the sense that they contain all information required to run the job. So the Accelerator job dispatcher actually just creates processes and points them to the job directory. New processes have to go and figure out their purpose by themselves by looking in this directory.

A running job has the process' *current working directory (CWD)* pointing into the job directory, so any files created by the job (including return values) will by default be stored in the job's directory.

When a job completes, the meta data files are updated with profiling information, such as execution time spent in single and parallel processing modes.

All code that is directly related to the job is also stored in the job directory in a compressed archive. This archive is typically limited to the method's source, but the code may have manually added dependencies to any other files, and in that case these will be added too. This way, source code and results are always connected and conveniently stored in the same directory for future reference.

3. Unique jobs are only executed once.

Among the meta information stored in the job directory is a hash digest of the method's source code (including manually added dependencies). This hash, together with the input parameters, is used to figure out if a result could be re-used instead of re-computed. This brings a number of attractive advantages.

4. Jobs may link to each other using job references.

Which means that jobs may share results and parameters with each other.

5. Jobs are stored in workdirs.

6. There may be any number of workdirs.

This adds a layer of "physical separation". All jobs relating to importing a set of data may be stored in one workdir, perhaps named `import`, and development work may be stored in a workdir `dev`, etc.

7. Jobids are created by appending a counter to the workdir name, so a job `dev-42` may access data in `import-37`, and so on, which helps manual inspection.

8. Jobs may dispatch other jobs.

It is perfectly fine for a job to dispatch any number of new jobs, and these jobs are called *subjobs*. A maximum allowed recursion depth is defined to avoid infinite recursion.

Chapter 5

Datasets

DRAFT

The *dataset* is the preferred way to store row-column data using the Accelerator. The dataset provides fast parallel streaming access to data, and strict typing with many supported types. Furthermore, datasets are lightweight – adding new columns to a dataset, or appending datasets to each other are instantaneous operations. Using datasets in a program is simple, since all operations are available through the `Dataset` class.

Datasets are created by methods, and are therefore located inside job directories. There can be any number of datasets in a job. The most obvious way to generate a dataset is using the `csvimport` method that creates a dataset from an input file. But the fact that a job can create and hold many datasets opens up for a lot of possibilities. The `csvimport` method, for example, creates a dataset from the input data, and it can be instructed to put unparseable lines of input data into a second “bad” dataset.

For performance reasons, datasets are split into several slices, and each data row exists in exactly one of the slices. The actual slicing may be carried out in different ways, like round robin, or randomly, but an interesting approach is to slice according to the hash value of a certain column. Slicing according to a hashed column ensures that all rows with a certain column value always ends up in the same slice. Hash-based slicing often makes completely parallel processing of a dataset possible, since related data is not spread over different slices, and thus no merge or “reduce” stage is required.

5.1 Dataset Internals

On a high level, the dataset stores a *matrix* of rows and columns. Each column is represented by a column name, or *label*, and all columns have the same number of rows. Columns are typed, and there is a wide range of types available. Typing will be introduced in section 5.8.

The dataset is further split into disjoint slices, where each slice holds a unique subset of the dataset’s rows. Slicing makes simple but efficient parallel processing possible. See Figure 5.1. The number of slices is set initially by the user in the Accelerator’s configuration file, and all workdirs that are used together in a project must use the same number of slices.

On a low level, there is one file stored on disk for each slice and column. A job that needs to read only a subset of the total number of columns may open and read from the relevant files only.

A technical note: If the number of slices is large and files are small, there will be a significant overhead from disk `seek()`, especially if using rotating disks. The Accelerator mitigates this by changing the storage model to using single files with offset-indexing when appropriate.

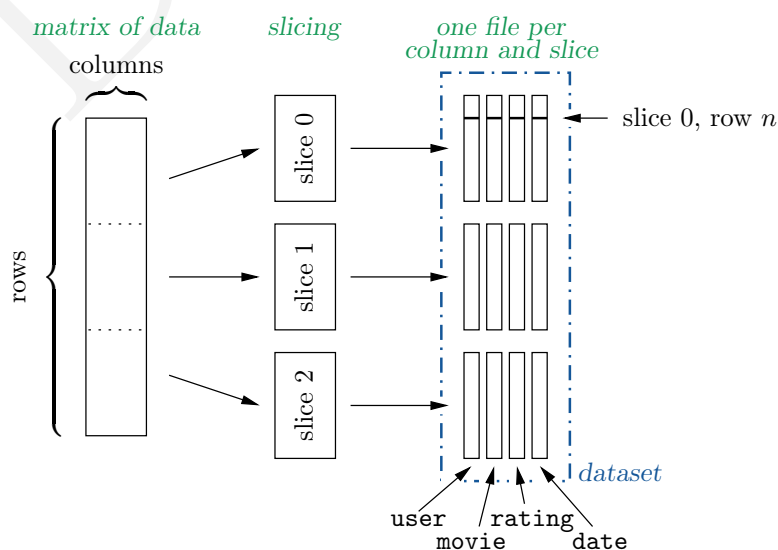


Figure 5.1: A “movie rating” dataset composed of four columns sliced into three slices.

5.2 Slicing and Hashing

As shown in the previous section and in figure 5.1, datasets are *sliced* into disjoint sets, called *slices*, that can be accessed independently. Typically, in a dataset, there is one file on disk for each slice *and* column. The main reason for doing this is performance. All files could be read in parallel, and only files relevant to the task at hand are read.

How to slice a dataset is not unique. There are many ways to carry out slicing. Perhaps the most simple way is to use round-robin, which cycles through the slices when writing, one data item at a time. Round-robin will balance the number of rows per slice as equal as possible, which is a good thing in many scenarios. In semi-mathematical terms, round robin would be

$$n \rightarrow n \bmod N$$

Meaning that input data row n is stored in slice $n \bmod N$.

Another way is to slice by looking at the values of a fixed single column and put all rows with equal “property” in the same column. This way, data will be sliced “by content”, and the number of rows per slice may vary significantly. In the context of the Accelerator, the “property” is a hash function, and the process is called *hash partitioning*. A dataset can be hashed on any single column, as long as its contents is hashable. Written as an equation, it will look like this

$$n \rightarrow \text{hash}(\text{data}) \bmod N$$

where “data” is the value of row n in the hashing column.

In many practical applications, data may be sliced using a hashing function so that data in each slice becomes *independent* for the application at hand. Independent data means that processing of the dataset can be carried out in a completely parallel fashion.

The hash function used by the Accelerator is a well-known function called siphash-2-4 that is available from the Accelerator’s `gzutil` library

```
from accelerator.gzutil import siphash24
```

```
y = siphash24(x)
```

This function is normally only used “under the hood”, there should be no need to call it explicitly.

5.3 Chaining

When a dataset is created, it is optional to attach a link to another dataset using the parameter `previous`. This is called *chaining*. Chaining provides a lightweight way to append rows to datasets, simply by linking datasets together. A typical use case is the import of log files. A new dataset is created from each new log file, and each dataset chains to the previous. Reading the full chain will access all log rows. As will be discussed in detail later, this relates to dataset *iterators* (see chapter 6), that may continue iterating over the next dataset in the chain when the current dataset is exhausted. Here is an example of how `csvimport` jobs can be chained

```
job = urd.build('csvimport', filename='file1.txt')
job = urd.build('csvimport', filename='file2.txt', previous=job)
```

In order to maintain high speed when processing long chains, the Accelerator caches chain metadata every 64th dataset. This reduces seek times significantly on rotating disks.

5.4 Dataset as Job Input Parameter

Datasets may be input to a method using the `datasets` input parameter list. In a running job, the items in this list are object of the `Dataset` class. See the following example

```

datasets = ('source',)
def synthesis():
    print(datasets.source.columns)

```

Note the comma used to indicate that this is in fact a Python `tuple` of dataset(s).

It is also possible to input a list or set of datasets, like this

```

datasets = (['source'],)
def synthesis():
    for ds in datasets.source:
        print(ds.columns)

```

5.5 Datasets from Job Objects

Information about a job instance’s datasets is provided using the `job.dataset()` function and the `job.datasets` attribute. To find all datasets in a job, use `job.datasets` like in this example

```

def main(urd):
    job = urd.build('create_datasets')
    for ds in job.datasets:
        print(ds.name)

```

To work on a specific dataset, just ask for it using its name as input parameter to `job.dataset()`,

```

ds_first = job.dataset('first')
ds_default = job.dataset()

```

Without an input parameter, the default dataset is returned. An error is issued if the dataset does not exist.

5.6 Dataset Properties

The `Dataset` class has a number of member functions and attributes that is intended to make it simple to work with. These functions will be described in the next sections. But first a note on naming datasets.

5.6.1 Dataset Name

The name of a dataset is accessible using the `.name` attribute, like this

```

print(ds_first.name)

```

The Accelerator is designed to handle various string encodings with ease, and in most situations the naming rules are very liberal. The dataset name, however, should preferably be *limited to ASCII characters*, since a directory with the name of the dataset will be stored on disk. The Accelerator cannot guarantee that the file system in use handles any “special” characters. Newline is for example not allowed.

5.6.2 Column Names

All columns in a dataset may be acquired using the `.columns` property, like this

```

datasets = ('source',)

def synthesis():
    print(datasets.source.columns.keys())
    # may print something like
    # ['GTIN', 'date', 'locale', 'subsource']

```

The `.columns` attribute is actually a dictionary from column name to properties, as will be shown in the next section.

Not all column names are valid, see section 5.9.7 for more information.

5.6.3 Column Properties

For each column, the name, type, and if applicable, the minimum and maximum values are accessible like this

```
print(datasets.source.columns['locale'].type)
# number

print(datasets.source.columns['locale'].name)
# locale

print(datasets.source.columns['locale'].min)
# 3

print(datasets.source.columns['locale'].max)
# 107
```

Creation of the `max` and `min` values is a simple operation that is done in linear time when the dataset is created. Maximum and minimum values are used for example when iterating over chains of sorted datasets, to quickly decide if a dataset is outside range and can be skipped in its entirety, see section 6.4.

5.6.4 Rows per Slice

It may be interesting to see how many rows there are per slice in a dataset. This information is available as a list, for example

```
print(datasets.source.lines)
# [5771, 6939, 6212, 6312, 6702, 6341, 5988, 6195,
# 6741, 6587, 6518, 5840, 6327, 5933, 6745, 6673,
# 6536, 6405, 6259, 6455, 6036, 6088, 6937, 6245,
# 6418, 6437, 6360, 6106, 6878]
```

The first item in the list is the number of rows in slice 0, and so fourth. The total number of rows in the dataset is the sum of these numbers.

5.6.5 Dataset Shape

The shape of the dataset, i.e. the number of rows and columns, is available from the `shape` attribute

```
print(datasets.source.shape)
# (4, 184984)
```

The second number is exactly the sum of the number of lines for each slice from above.

5.6.6 Hashlabel

If the dataset is hash partitioned on a particular column, the name of this column is stored in the `hashlabel` attribute

```
print(datasets.source.hashlabel)
# GTIN
```

5.6.7 Filename and Caption

The dataset may have a filename associated to it. This makes sense in situations for example where the dataset is created from an input data file using `csvimport` or similar. The filename is accessible using the `filename` attribute:

```
print(datasets.source.filename)
# /data/incoming/raw_repository_5391.gz
```

Furthermore, it is possible to set a caption at dataset creation time. The caption is entirely user-defined and has no function in the Accelerator. The caption is accessible like this

```
print(datasets.source.caption)
# hash_partitioned_data
```

5.7 Operations on Chains

The `.chain` function is used to operate on dataset chains. It takes a dataset as input, some options that will be discussed next, and returns a `DatasetChain` object.

```
# return a chain object for dataset "ds"
chain = ds.chain()
```

The `chain()`-function takes the following optional arguments

name	default	description
<code>length</code>	<code>-1</code>	Number of datasets to include, default is <code>-1</code> , meaning all datasets in chain. Do not confuse with <code>stop_ds</code> .
<code>reverse</code>	<code>False</code>	Return the datasets in reverse order. Default <code>False</code> .
<code>stop_ds</code>	<code>None</code>	Return datasets <i>from</i> <code>stop_ds</code> to current dataset. Do not mix with <code>length</code> .

The returned value, `chain` in the previous example, is of type `DatasetChain`, which supports the following operations.

name	description
<code>min(<column>)</code>	Minimum value of column, or <code>None</code> if column does not exist or is not sortable.
<code>max(<column>)</code>	Return minimum value of column, or <code>None</code> if column does not exist or is not sortable.
<code>lines(sliceno=<code>None</code>)</code>	Default is number of lines in chain. With option <code>sliceno=x</code> , number of lines in slice <code>x</code> .
<code>column_count(<column>)</code>	Number of datasets in chain containing column <code><column></code>
<code>column_counts()</code>	Counter of occurrences per column per dataset in chain.
<code>with_column(<column>)</code>	A <code>DatasetChain</code> object containing only those datasets that has column <code><column></code> .
<code>iterate(...)</code>	Same arguments as <code>Dataset.iterate()</code> . Will iterate over the whole chain.

5.8 Column Data Types

Dataset columns are typed. This means, for example, that if a column's type is `date`, each value read from the column will be in Python's `date` format, ready for processing. The same goes for all supported types, including `json` and `pickle`, that may return rather complex datatypes.

By default, a typed column does not allow the storage of `None` values. This can be changed by setting the `none_support` boolean when creating the column, see section C.6.1.

All available types are shown in the following table. More details follow in the next sections.

name	description
number	float or int ($\pm(2^{1007} - 1)$)
float64	64 bit float (this is what the Python <code>float</code> type uses)
float32	32 bit float (this has lower precision)
int64	64 bit int ($\pm 2^{63} = \pm 9223372036854775807$)
int32	32 bit int ($\pm 2^{31} = \pm 2147483647$)
complex64	64 bit complex number (i.e. two float64)
complex32	32 bit complex number (i.e. two float32)
bool	<code>True</code> or <code>False</code>
date	<code>datetime.date</code>
time	<code>datetime.time</code>
datetime	<code>datetime.datetime</code>
bytes	raw data
ascii	string limited to ascii
unicode	normal string
json	anything jsonable (anything <code>json.dump</code> accepts)
pickle	anything picklable (anything <code>pickle.dump</code> accepts)

Additionally a number of types prefixed with `parsed:` are available. These take everything the non-parsed version takes, plus strings that can be parsed as that type. Once written they are the type after the `parsed:` prefix.

name	notes
<code>parsed:number</code>	
<code>parsed:float64</code>	
<code>parsed:float32</code>	
<code>parsed:complex64</code>	strings use <code>j</code> , not <code>i</code>
<code>parsed:complex32</code>	strings use <code>j</code> , not <code>i</code>
<code>parsed:int64</code>	also accepts floats
<code>parsed:int32</code>	also accepts floats
<code>parsed:json</code>	strings are always decoded, even though they would of course be accepted as-is by the plain <code>json</code> type

5.8.1 Arbitrary precision numbers: number

The type `number` is integer when possible and float otherwise. It can handle very large numbers, up to $\pm(2^{1007} - 1)$. This is usually slower than the fixed size types, but not by much.

5.8.2 Standard Fixed Size Numbers

The common `int` and `float` types in 32 and 64 bit versions are available for use when the range of the data is known.

Note that the `float32` type not only has a smaller range than `float64`, but also lower precision. The python `float` type has the same precision as `float64`, so that is the safe choice. If you use `float32` you may be surprised by e.g. `0.1` not being the same value after you read it back.

5.8.3 Booleans

The `bool` type is used to store logical `True` or `False` values only.

5.8.4 Types Relating to Time

The `date`, `time`, and `datetime` types use the corresponding classes from the standard `datetime` module. A column that is typed to any of these may directly take advantage of the high level time related methods, like for example

```
for ts in datasets.source.iterate(sliceno, 'timestamp'):
    print(ts.strftime('%Y-%m-%d '))
```

5.8.5 String Types

There are two string types, `unicode` and `ascii`. Normally you want to use the `unicode` type, which allows all valid unicode strings (but not surrogates, `\ud800` to `\udfff`, which python can put in strings decoded with the `'surrogateescape'` error handler).

Sometimes the more restrictive `ascii` type, which only allows `\u0000` to `\u007f` can provide a welcome restriction.

On Python2, the `ascii` type will return `str` objects (a.k.a. `bytes`) and `unicode` will return `unicode` objects.

On Python3, both return `str` objects (i.e. what Python2 called `unicode`.)

5.8.6 Raw Data

The `bytes` type is used to store raw data, such as binary image files. The upper storage limit for a value typed as `bytes` is almost 2GB ($2^{31} - 1$ bytes). The `csvimport` standard method uses this type for all data in its output dataset.

5.8.7 Complex Numbers

There are two complex number types, `complex32` and `complex64`. The `complex32` type has the same caveats as the `float32` type.

5.8.8 JSON

The `JSON` type makes it possible to store and load more complex data structures in a dataset. Anything that is `JSONable` works as input. Conversion between `JSON` and Python data types is done by the writers and iterators, so the user can just work on the data and never has to see the actual `JSON`, for example

```
dw = job.datasetwriter(...)
...
a = dict(x=3, y=dict(z=5, w=[1,2,3]))
dw.write(a)
```

5.8.9 Pickle

The Python `pickle` type is the most flexible type of all. It supports any `pickle`-able object, and encoding and decoding is done on-the-fly under the hood. This means that *it is possible to write almost anything into a dataset column!*

This type is however not support on Python2.

5.8.10 None-Handling

The value `None` is valid input for all types if you enable `none_support`, and `json` and `pickle` accept them regardless. For example, valid values for a `bool` type column are `{True, False}` without `none_support`, and `{True, False, None}` with `none_support`. See section C.6.1 for more information on `none_support`.

5.8.11 Does Hash Partitioning Depend on the Data Type?

Hash partitioning is independent of type where possible. The hashing function is designed so that, for example, the floating point value 4.0 has the same hash value as the integer value 4. So for integers, the result is the same independent of floating point of integer representation. But note that hashing is different for strings, the hash of the string "4" is *not* the same as the hash of the integer 4.

5.9 Creating a New Dataset

Datasets are created by methods using the `DatasetWriter` class. A factory for this class is available in a running method as `job.datasetwriter` like this

```
def prepare(job):  
    dw = job.datasetwriter()
```

The most common scenario is to set up the new dataset in `prepare()`, and write data to it in parallel in `analysis()`, but it is also possible to write a dataset in an entirely serial fashion in `prepare()` or `synthesis()`. When a dataset-creating method terminates, it will create and store all required meta-information, such as min/max values, for the created dataset(s) automatically.

The most common arguments to `DatasetWriter` are

name	description
filename	if there is a filename associated, store it here
caption	additional caption
hashlabel	name of column to hash by when slicing
previous	previous Dataset, for chaining
name	dataset name, default set to <code>default</code>
parent	parent Dataset when adding columns

5.9.1 Create in `prepare()` + `analysis()`

The following example will use `DatasetWriter` to create a dataset with three columns of different types. The name of the dataset will be `firstset`. The writer will be initialised in `prepare()`, and data will be written to the dataset in `analysis()`. Note that the example actually creates a dataset *chain*, since it links the dataset under creation to the dataset named `previous` from the input parameters.

```
datasets = ('previous',)  
  
def prepare(job):  
    dw = job.datasetwriter(  
        previous=datasets.previous,  
        name='firstset',  
    )  
    dw.add('X', 'number')  
    dw.add('Y', 'unicode')  
    dw.add('Z', 'time')  
    return dw  
  
def analysis(sliceno, prepare_res):  
    dw = prepare_res  
    ...  
    for x, y, z in some_data:  
        dw.write(x, y, z)
```

The function `dw.write()` is used to write data to the dataset. The order of the variables in the `.write()` function call is the same as the order of the `.add()` calls in `prepare`. There are a few alternative ways of writing data, as shown here

```
# write a dict with keys corresponding to column names
dw.write_dict({column: value})

# write a list with items in same order as dw.add() calls
dw.write_list([value, value, ...])

# one parameter for each .add() call, in same order
dw.write(value, value, ...)
```

Several datasets can be created simultaneously using multiple writers with different names.

5.9.2 Create in `synthesis()`

Since a dataset is sliced in multiple disjoint sets, and `synthesis()` is run only once, data has to be sliced during writing somehow. There are two possible ways to do this. One is to first set a slice number

```
dw.set_slice(sliceno)
```

before writing data into that slice. Note that this can only be done once per slice. The other is to use one of the `split_write` functions

```
# use a dict-writer
writer = dw.get_split_write_dict()
writer({column: value})

# use a list-writer
writer = dw.get_split_write_list()
writer([value, value, ...])

# use a parameterised writer
writer = dw.get_split_write()
writer(value, value, ...)
```

These writers will write round-robin if the dataset is not hashed, and to the “matching” slice if the dataset is hash partitioned.

5.9.3 Creating a Dataset Chain in one Job

It is possible to create a chain of datasets by inserting a writer object as previous to another writer, like this.

```
def prepare(job):
    dw1 = job.datasetwriter(name='ds1')
    dw2 = job.datasetwriter(name='ds2', previous=dw1)
```

5.9.4 Completing Dataset Creation

Normally, there is no need to tell the `DatasetWriter` that the last line of data is written. This is handled automatically when the method exits. In some situations, such as when a dataset is to be used by a `subjob` launched from the creating method, it is necessary to manually tell the writer that the dataset is complete. This is done by calling `finish()` as shown below

```
dw.finish()
```

The `finish()`-call returns a dataset object, so the just finished dataset could be put in use immediately, like this

```
ds = dw.finish()
it = ds.iterate(None, 'user')
```

5.9.5 Datasets Created by Subjobs

If a dataset is created in a subjob, it is not visible from the build scripts. This is solved by linking dataset meta-information to the job calling the subjob, using the `.link_to_here()` function. This is explained in detail in section 4.10, but the basic idea is as follows

```
from accelerator import subjobs
def synthesis():
    job = subjobs.build('dataset_creator')
    ds = job.dataset('ds1')
    ds.link_to_here('first_dataset')
```

Here, the subjob creates a dataset named `ds1`. This dataset resides in the subjob’s job directory, but a link to it from the *current* job is created using the `link_to_here` function, which also allows the dataset to be renamed. The current job will now claim that it contains a dataset `first_dataset`.

In the same call, the subjob’s dataset’s previous can also be overridden:

```
ds.link_to_here('first_dataset', override_previous=anotherds)
```

This might sound complicated, but it simply means that the current job can export a set of datasets created by its subjobs, and that it can manipulate these datasets so that they may appear to be chained in any way the user finds fit.

5.9.6 Creating Hash Partitioned Datasets

A hash partitioned dataset is created by setting the `hashlabel` argument to `job.datasetwriter`. For a hash partitioned dataset, only data fulfilling the hashing requirement for a slice may be written to that slice, and an exception will be raised if the data to be written does not belong to the current slice.

A simple way to filter the data to be written is to call

```
dw.enable_hash_discard()
```

first in each slice or after each `.set_slice()`. Then, writes that belongs to another slice are silently ignored, while “correct” data gets written as expected.

It is possible to check before writing if the data is to be put into the current slice using the `dw.hashcheck()` function, like this

```
...
if dw.hashcheck(hashcoldata):
    # compute bulkdata here
    bulkdata = expensive_function(...)
    # and write to dataset
    dw.write(hashcoldata, bulkdata)
```

This is beneficial if it is expensive to compute the data to be stored. In the example above using `hashcheck()`, data is only computed if it is to be stored in the slice.

In general, the hash function for a particular column is available like this

```
dw.writers[colname].hash
```

This function can be used to manually check if the data belongs to a slice. For more details and alternatives, please see the documentation in the source file `dataset.py`.

5.9.7 Column Name Restrictions

Column names must be valid Python identifiers. Invalid characters are replaced by the underline (`_`) character. The underline character is also used to make column names unique when necessary. The table below shows some examples.

input	converted	comment
"_"	"_"	Converting to valid python identifier.
"a b"	"a_b"	Converting to valid python identifier.
"42"	"_42"	Converting to valid python identifier.
"print"	"print_"	print is a keyword (in py2).
"print@"	"print__"	print_ was just taken.
"None"	"None_"	None is a keyword (in py3).

5.9.8 More Advanced Dataset Creation

Currently out-of-scope of this manual. Please see the source file `dataset.py` for full information.

5.10 Appending New Columns to an Existing Dataset

With minimal overhead, existing datasets could be extended with new columns. Internally, this is implemented by storing the new column data together with a pointer to the original, “parent”, dataset.

Appending new columns works the same way as when creating a dataset, with the exception that a link to a dataset that is to be appended to is input to the writer constructor. Columns can be appended either in `analysis` or `synthesis`, as shown in the two following sections. Note that appending a column does only apply to one single dataset, and not to the complete chain of datasets, if present. See the blog at exax.org for several examples of how to append columns to a dataset chain.

5.10.1 Appending New Columns in Analysis

The following example appends one column to an existing dataset `source`, while chaining to the dataset `previous`.

```
datasets = ('source', 'previous',)

def prepare(job):
    dw = job.datasetwriter(
        parent=datasets.source,
        previous=datasets.previous,
        caption='with the new column',
    )
    dw.add('newcolname', 'unicode')
    return dw

def analysis(sliceno, prepare_res):
    dw = prepare_res
    for data in datasets.source.iterate(sliceno, ... )
        ...
        dw.write(value)
```

The `DatasetWriter` will automatically check that the number of appended rows does match the number of rows in the parent dataset. Otherwise, an error will be issued and execution will terminate. Typically new columns are derived from existing ones, so this is usually not a problem.

5.10.2 Appending New Columns in Synthesis

A straightforward way to append columns to a dataset in synthesis is using the `set_slice()` function, as shown in the example below.

```

def synthesis(job):

    dw = job.datasetwriter(parent=datasets.source)
    dw.add('newcolumn', 'json')

    for sliceno in range(job.params.slices):
        dw.set_slice(sliceno)
        for data in datasets.source.iterate(sliceno, ...):
            ...
            dw.write(x)

```

Note that the `for`-loop over all slices is controlling *both* the reading iterator *and* the dataset writer. Note also that `.set_slice()` can only be called once per slice.

5.11 Standard Methods on Datasets

Figure 5.2 shows an example of how the standard methods `csvimport`, `csvexport`, `dataset_sort`, `dataset_hashpart`, and `dataset_unroundrobin` can be used to change the storage of a dataset.

In the upper part of the figure, a file `data.csv` in a column storage format is imported using the `csvimport` method. This method reads one line at a time and writes to the a new dataset’s slices in a round robin fashion.

In many applications, this round robin reorder does not matter, but in some cases, such as when the input data is a time series, it is critical to keep the order of the data intact. In those situations, the `dataset_unroundrobin` is be used to “transpose” the dataset. One can think of is as if the method reads data row by row, but writes it slice by slice. If the `csvexport` method is applied to an “unroundrobined” dataset, the output CSV file will be in the same order as the original input file.

The `dataset_unroundrobin` method takes a parameter `trigger_column` which is used to make sure that a slice switch does not appear as long as the data in that column is a constant value. If the data is sorted, this ensures that all rows with a particular value is not spread along several slices. In figure 5.2, the transition from slice 1 to slice 2 appears while the method is writing the value “10”. If `trigger_column` is enabled, switching is delayed until all “10”s are written.

The `dataset_hashpart` method is used to partition the data into disjoint slices based on the content of a single column. Depending on the distribution of the data, the generated dataset will be more or less balanced.

Finally, the `dataset_sort` method is used to sort data. It can sort data independently per slice, or sort all the data in a dataset across all slices. In the latter case, there is a `trigger_column` option similar to the one in `dataset_unroundrobin` that ensures that slice switches appear only when the value of a particular column is changing. This is the reason for the first column in the bottom right dataset in figure 5.2 to be so long. Slice switching should happen while writing the “10” value, but since the `trigger_column` is set, switching is delayed until all “10”s are written.

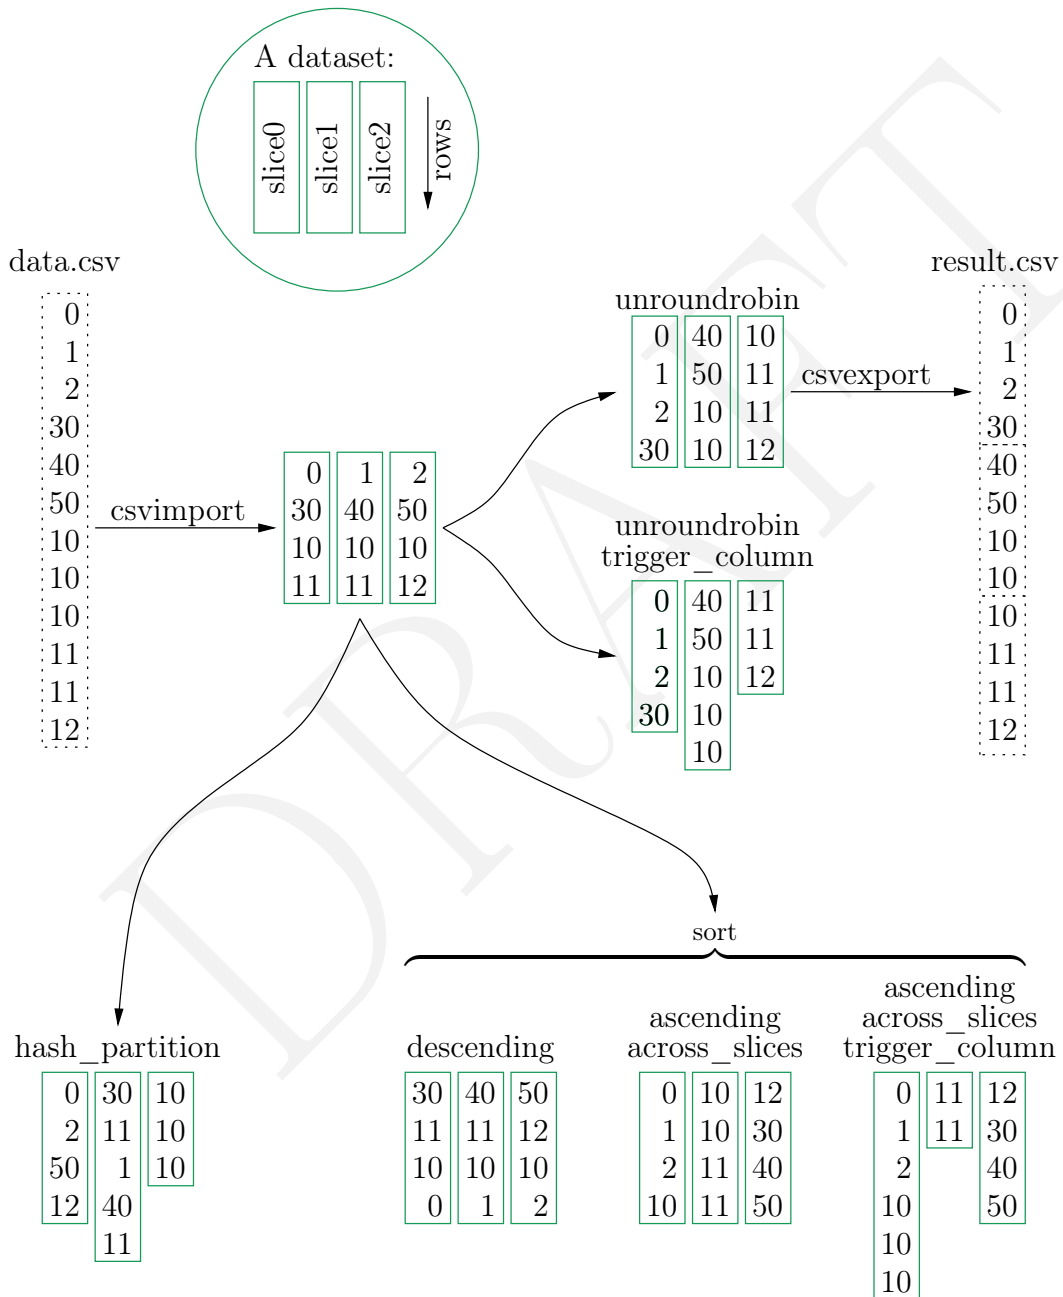


Figure 5.2: An example of how a dataset (with three slices) is modified by some common standard methods.

Chapter 6

Data Iterators

DRAFT

The basic idea of the Accelerator's datasets is to make it easy to create parallel programs that can read and write large amounts of data at a very high speed. High speed data read access is implemented as a set of special Python *iterators*. Each iterator yields one `tuple` at a time containing elements from one or more specified data columns, one row at a time. In case of iterating over a single column, the output may optionally be a scalar instead of `tuple` for cleaner code and more efficient computing.

6.1 The Three Iterators

Technically, iterators are members of the `Dataset` class. Iterators can be parallel, in `analysis()`, or sequential, in `prepare()` or `synthesis()`. There are three iterators available:

```
iterate(), for single dataset iteration,  
iterate_chain() for iterating over dataset chains, and  
iterate_list() for iterating over a specified list of datasets.
```

And each of them will be discussed later in this chapter. For completeness, it should also be mentioned that the `DatasetChain` class also has an `.iterate()` function, and it works similar to `iterate_chain()`.

In many common use cases it is sufficient to provide only two arguments to the iterator: `sliceno`, which is mandatory, and `columns`. These, and all other arguments are presented in detail shortly. A typical use of an iterator looks like this

```
datasets = ('source',)  
  
def analysis(sliceno):  
    for m, u in dataset.source.iterate(sliceno, ('movie', 'user')):  
        # do something with m and u here...
```

Python's constructors can be used to create objects from iterators like in the following example, where the purpose is to compose a dict.

```
n2d = dict(dataset.source.iterate(sliceno, ('name', 'date',)))
```

or this example

```
from collections import Counter  
...  
c = Counter(dataset.source.iterate(sliceno, 'user'))
```

that will create a counter of how many times each user is present in the dataset.

6.1.1 Iterator Arguments

All three iterators share these arguments

name	default	description
<code>sliceno</code>	<code>mandatory</code>	Slice number (an integer) to iterate over, <code>None</code> to iterate over all slices sequentially, or <code>roundrobin</code> to take one value per slice in a round robin fashion.
<code>columns</code>	<code>None</code>	Tuple of column labels or a single name if iterating over one column. <code>None</code> selects all columns in alphabetical order.
<code>hashlabel</code>	<code>None</code>	Name of hash column. If the code relies on a dataset being hashed on a particular column, set this to make the iterator <code>assert</code> that this is actually the case. Execution will terminate if the <code>hashlabel</code> is incorrect.
<code>rehash</code>	<code>False</code>	Setting this to <code>True</code> will hash partition the dataset on-the-fly based on the <code>hashlabel</code> column. (Rehashing on-the-fly is slower, so ideally datasets should be rehashed using the <code>dataset_rehash</code> method 8.5.)
<code>status_reporting</code>	<code>True</code>	Give status when pressing C-t. Unless manually zipping several iterators together, this should be set to default <code>True</code> . See <code>dataset.py</code> source code for full information.

In addition, `iterate_chain` takes these arguments too

name	default	description
<code>length</code>	<code>-1</code>	Number of datasets in a chain to iterate over. Default is <code>-1</code> , which corresponds to all datasets in a chain.
<code>range</code>	<code>None</code>	Filter rows based on a column's value being within a range, see section 6.4
<code>sloppy_range</code>	<code>False</code>	Used with <code>range</code> , but will iterate over full datasets for those datasets in a chain that have values within range, see section 6.4.
<code>reverse</code>	<code>False</code>	Iterate chain backwards. Default is to iterate forward, i.e. from oldest to newest dataset.
<code>stop_ds</code>	<code>None</code>	Iterate back to this dataset. Actually, setting this will iterate from the dataset <i>following</i> <code>stop_ds</code> to the newest dataset in the chain.
<code>pre_callback</code>	<code>None</code>	A function that will be called before iterating each dataset.
<code>post_callback</code>	<code>None</code>	A function that will be called after iterating each dataset.

and `iterate_list()` takes a `datasets` parameter

name	default	description
<code>datasets</code>	<code>None</code>	List of datasets to iterate over.

6.2 Basic Iteration

Basic use include iterating in parallel or serial over one dataset or a chain of datasets.

6.2.1 Parallel Iterator Invocation

For parallel iteration in `analysis()`, the iterator needs to know the number of the current slice. This information can be fed to the `analysis()` function in the `sliceno` variable. The following is an example of iteration that happens independently in each slice.

```
from collections import defaultdict
datasets = ('source',)

def analysis(sliceno):
    h = defaultdict(set)
    for user, item in datasets.source.iterate(sliceno, columns=('user', 'item')):
        h[user].add(item)
```

The program creates dictionaries mapping users to sets of items for the `source` dataset. Assuming that the dataset is hash partitioned (see 5.2), this operation is entirely parallel and there is no need to merge all the results from the analysis processes afterwards, since the different slices do not share any keys with each other.

6.2.2 Sequential Iterator Invocation

Setting the `sliceno` parameter to `None` will cause the iterator to run through all slices of the dataset, one slice at a time, like in this example

```
def synthesis():
    h = defaultdict(set)
    for user, item in datasets.source.iterate(None, columns=('user', 'item')):
        h[user].add(item)
```

Dataset slices will be iterated in increasing order.

6.2.3 Special Case, Round Robin Iteration

By default, the iterators stream slices of data. This is almost always exactly what is needed. But sometimes, for example when the order of rows imported by `csvimport` matters, there is a need for a row-wise iteration order. For maximum performance, the `csvimport` method writes datasets in a round robin fashion, so iterating over a `csvimported` dataset does not return the lines in the same order as they were written.

By setting the first parameter of any of the iterator functions to “`roundrobin`”, the iterator will internally fetch all slice iterators and return one value at a time from each iterator in a round robin fashion. The resulting output is then in the same order as in the file imported by `csvimport`. In a dataset chain, round robin will happen *per dataset*. There is a performance penalty associated with this functionality, but it is handy for time-series-like data.

6.2.4 Special Cases, Iterating Over All or a Single Column

It is possible to iterate over all columns in a dataset by specifying an empty list of column names, like this

```
for items in dataset.source.iterate(sliceno, None):
    print(items)  # is a tuple of all columns
```

The iterator will output a `tuple` populated with all column values for each row. The columns will be in sorted column name order.

If iterating over a single column, it makes little sense to keep the output values in a one-dimensional tuple. A scalar is cleaner and more efficient. Here are the two different ways to iterate over a single column

```

# alternative 1, use lists/tuples
for user in datasets.source.iterate(sliceno, ('USER',)):
    userset.add(user[0]) # user is a tuple

# alternative 2, specify column as string, not list
for user in datasets.source.iterate(sliceno, 'USER',):
    userset.add(user) # user is a scalar!

```

6.2.5 Iterate Over Chains

The `iterate_chain()` iterator is used to iterate over one or more datasets in a chain, starting at the “oldest” dataset. The following example will iterate over the last three datasets in the chain, oldest dataset first.

```

datasets = ('source',)

def analysis(sliceno):
    h = defaultdict(set)
    for user, item in datasets.source.iterate_chain(
        sliceno, columns=('user', 'item',), length=3,
    ):
        h[user].add(item)

```

Using `iterate_chain()` without explicitly specifying `length` will default to a `length` of `-1`, which corresponds to all datasets in the chain.

6.3 Halting Iteration

Iteration over a dataset chain will continue until all datasets are exhausted or a stop criteria is fulfilled. There are several mechanisms for stopping, and they may be combined in a single iterator call. If so, iteration will be over the shortest range of the conditions.

6.3.1 Halting Using `length`

```

for user, item in datasets.source.iterate_chain(
    sliceno,
    ('user', 'item',),
    length=options.length,
):

```

This will iterate for the last `options.length` number of datasets. Note that a `length` of `-1` is default and will iterate without bounds.

6.3.2 Halting Using `stop_ds`

Similar to using `length`, but will stop when reaching a certain dataset.

```

for user, item in datasets.source.iterate_chain(
    sliceno,
    ('user', 'item',),
    stop_ds='foo-3',
):

```

Stopping at a constant dataset has limited value. Next section shows how to stop iterating based on previous jobs.

6.3.3 Halting Using Another Job’s Input Parameters

```

for user, item in datasets.source.iterate_chain(
    sliceno,
    ('user', 'item',),

```

```

    stop_ds={jobs.previous: 'source'},
):

```

This will iterate until reaching the source dataset of the `jobs.previous` job.

6.4 Iterating Over a Data Range

It is possible to iterate over rows having a specified column's value within a certain range. This works best on datasets that are sorted on the specified column.

```

for user, item in datasets.source.iterate_chain(
    sliceno,
    ('user', 'item',),
    range={'ts_column': (datetime(2016, 1, 1), datetime(2016, 3, 31))},
):

```

This example will limit the iterator to exactly the range of lines that fulfill the range condition. It is relatively costly to filter each line, and there is a speed advantage to specifying `sloppy_range=True`, which will iterate over all datasets that contain part of the range:

```

for user, item in datasets.source.iterate_chain(
    sliceno,
    ('user', 'item',),
    range={'ts_column': (datetime(2016, 1, 1), datetime(2016, 3, 31))},
    sloppy_range=True,
):

```

Here, all datasets that *contain* any line containing values within the range will be included in the iteration. Still, if the datasets are sorted, and there are many datasets, the side-effect caused by reading too many lines will be limited.

6.5 Iterating in the Reverse Direction

By default, iterating over a chain of dataset starts at the oldest dataset and ends at the latest dataset. This behavior can be reversed by specifying `reverse=True`. But note that row iteration is still in the forward direction within each dataset!

```

for user, item in datasets.source.iterate_chain(
    sliceno,
    ('user', 'item',),
    reverse=True,
):

```

6.6 Hash Partitioned Datasets and on-the-fly Partitioning

Hash partitioning a dataset on a particular column, see section 5.2, may really simplify the parallel programming of methods using the dataset. However, the parallel code will not work properly if it turns out that the input data is in fact not hash partitioned in the expected way. For that reason, it is a good idea to *assert* the hashlabel by entering it into the iterator function, like this

```

s = dict(datasets.source.iterate_chain(
    sliceno, ('user', 'item',), hashlabel='user',
))

```

so that execution will terminate if the `hashlabel` is not correct.

It is possible to hash partition the dataset on-the-fly. This is done by setting the `rehash` argument to the iterator to `True`, like this

```

for user, item in datasets.source.iterate_chain(
    sliceno,
    ('user', 'item',),
    rehash='item',
):
    # only lines with items such that
    # has(item) % slices == sliceno here

```

While this works, the preferred way to rehash is to use the `dataset_rehash` method 8.5, since it will store the rehashed dataset for later use, which in most scenarios will be more efficient.

6.7 Callbacks

The iterator may be assigned callback functions that are called before starting iterating a new dataset, and/or after the current dataset is exhausted. Callbacks are useful for example to aggregate data by dataset when iterating over a large dataset chain.

There are two independent callbacks for these two cases, called `pre_callback` and `post_callback`. If `sliceno` is set to `None`, i.e. iteration runs over all slices of all datasets in one process, it is even possible to have callback between slice changes.

The example below will print the dataset identifier for each dataset prior to iterating over it.

```

# pre_callback once per dataset
def prefun(dataset):
    print(dataset.name)

for user, item in datasets.source.iterate_chain(
    sliceno,
    ('user', 'item',),
    pre_callback=prefun,
):
    ...

```

The argument to the callback is the dataset instance corresponding to the dataset to be iterated next.

Next is an example of an iterator running over all slices. The callback function is executed before each new slice is iterated. The callback takes two arguments in this scenario, first, the dataset instance as per the example above, and second the number of the slice.

```

# callback once per slice
def prefun(dataset, sliceno):
    print(dataset.name, sliceno)

for user, item in datasets.source.iterate(
    None,
    ('user', 'item',),
    pre_callback=prefun,
):
    ...

```

The `post_callback` function is defined similarly.

6.7.1 Skipping Datasets and Slices from Callbacks

It is possible to skip dataset iterations by raising exceptions, as follows.

- To skip the next dataset do

```
raise SkipJob
```

- To have the iterator skip a slice, do a

```
raise SkipSlice
```

- And to abort iterating completely

```
raise StopIteration
```

In this case, a `post_callback` will never be run.

DRAFT

Chapter 7

High Level Control: Urd

DRAFT

This chapter is the continuation of chapter 3, “Basic Build Scripting”. Please read about builds script and joblists before proceeding.

7.1 Introduction to Urd

Urd is a transaction log-based database and server used to group, tag, and keep track of jobs. Build scripts are capable of a large variety of complicated tasks, and on top of that Urd adds the capabilities of job organisation, storage, and retrieval. This makes it possible to handle much larger and more advanced projects while the process is fully transparent and reproducible.

Using Urd, a project can be separated into functionally independent parts, or *sessions*, each containing at least one job, and all dependencies between jobs inside as well as between these sessions are tracked by a *transaction log*. Sufficient information is stored so that any part or session of a project could theoretically be re-constructed the way it looked at any instance in time. More formally, Urd provides two things:

1. Separation between build scripts, and a way to share information about existing sessions between different scripts (or the same script at different points in time).
2. A searchable transaction log database of all sessions, together with their dependencies on other session. A timestamp, date, integer, or combination thereof can be used as key.

Due to its transaction log, there is no way using the provided API to accidentally modify or destroy data stored in the Urd database. The interesting transaction log database and many other aspects of Urd will be explained in this chapter.

7.2 A Simple Use Case

A simple use case is presented here to illustrate the basic Urd features. Assume a project where movie recommendation data is to be analysed. Every hour, new data corresponding to the last hour’s activity is added to the project in the shape of a log file. The project is based on two build scripts, one *import* script, and one *analysis* script.

7.2.1 The Import Build Script

The import scrip is used to look for new data files, and import and chain them as they become available. For each new file that is imported, the build script will tell the Urd server the *timestamp* of the file as well as a *list of all jobs*, i.e. the Urd session, that have been built based on that file. Typically, this includes `csvimport`, `dataset_type`, perhaps a `dataset_hash_partition`, and similar data pre-processing jobs.

7.2.2 The Analysis Build Script

The analysis build scrip is used for the data analysis work, and is perhaps run less regularly. This script needs to know which jobs that correspond for the most recently imported file, and this is a straightforward thing to ask Urd. When all processing is done, all the analysis jobs are also stored in Urd together with a corresponding timestamp so that they can be retrieved later.

7.2.3 Urd Provides Separation

In the example, Urd is used to forward information about executed jobs from the first build script (the import script) to the second (the analysis script). In this sense, Urd provides *isolation* by “message passing” between build scripts. The import scripts can be modified or even removed, but as long as no workdirs are cleared and the Urd database is intact, it is still possible to fetch references to the existing import jobs from Urd. Thus, it is always possible to go back and see what things looked like at any previous point in time.

7.2.4 Transparency

Urd also provides *transparency*, since it can tell which input files that were used by any analysis job. In this way, it will never happen that there is uncertainty about which data that was used for a particular analysis.

To investigate the result of an analysis, Urd can be asked to return the jobs corresponding to, say, the most recent analysis jobs. The query response from Urd will contain these jobs, *as well as* the results from any Urd queries that were carried out for the jobs to run. This information will contain the import jobs, since these were queried in order to perform the analysis.

7.3 Local or Shared Urd Server

By default, Urd is run as a *local* server, which means that it is not externally accessible. A non-local Urd server that can share information between several users is straightforward to set up, see section B.5.

7.4 Urd Sessions and Lists

A simple file import script will be used as example in this section:

```
def main(urd):
    urd.build('csvimport', filename='txn1.txt')
```

In order to use this import job in a future context, a *session* is created by wrapping the code by the `urd.begin()` and `urd.finish()` functions, like this

```
def main(urd):
    urd.begin('import/txn', '2018-05-03')
    urd.build('csvimport', filename='txn1.txt')
    urd.finish('import/txn')
```

Everything that happens between `.begin()` and `.finish()` makes up the session. The `.finish()` function makes sure that the session is stored permanently to disk for future reference. In this case, the session can be retrieved knowing its *list* name

```
import/txn
and its timestamp
2018-05-03
```

The list identifier is composed of two parts, `<user>/<list>`, where `<user>` is for authorisation purposes. Each user can have any number of lists that are globally readable, but only the authorised user can write to them, as will be explained later.

7.5 A First Urd Query

The list created in the previous section is stored in the Urd database and therefore ready to be used by other build scripts. For example, here is a build script that does some processing on the previously imported file

```
def main(urd):
    urd.begin('process/test')

    import_session = urd.latest('import/txn')

    import_timestamp = import_session.timestamp
    import_job       = import_session.joblist['csvimport']

    urd.build('process', source=import_job)
```

```
urd.finish('process/test', import_timestamp)
```

The first thing that happens is that all processing is covered in a session named `process/test`. The timestamp can be set in *either* `.begin()` or `.finish()`.

The script is then retrieving the most recently created urd session stored in list `import/txn`. Two things are extracted from this data, the *timestamp* and the *joblist*. The timestamp will be used for this session as well, to indicate that processing is based on data with that particular timestamp. A reference to the `csvimport` job is then extracted from the joblist and fed to the `process` job as an input dataset parameter. (For information about the `JobList` class, see section 3.2.)

7.6 The Contents of a Stored Session

Calling `urd.finish()` will update the Urd database with the contents of the current *session*. Each session is addressable using a *list* name (in the format `<user>/<list>`) and a *timestamp*. Session data is stored internally in the JSON format, and in build scripts it will appear as a Python `dict`. The example presented earlier in this chapter may have been recorded similarly to this

```
{
  "user": "processing",
  "automata": "test",
  "timestamp": "2018-05-03",
  "caption": "",
  "joblist": [
    [
      "process",
      "TEST-37"
    ]
  ],
  "deps": {
    "import/txn": {
      "timestamp": "2018-05-03",
      "caption": "",
      "joblist": [
        [
          "csvimport",
          "TEST-34"
        ]
      ]
    }
  }
}
```

(This example states that at timestamp 2018-05-03 in list `processing/test`, there exists a `process` job TEST-34 that used a `csvimport` job TEST-34. This job also exists in the urd list `import/txn` at timestamp 2018-05-03.)

The most important keys are

name	description
timestamp	Timestamp of session
caption	A caption
user/automata	Name of Urd list
joblist	An object of type <code>JobList</code> , containing all jobs built in the session. For more information, see 3.2.

deps A dictionary of dependencies from user/automata to urd sessions: {'user/automata': session}.

7.7 Urd Sessions: .begin() and .finish()

There are a number of options associated with a session, as shown here,

```
urd.begin(urdlist, timestamp, caption=None, update=False)
urd.finish(urdlist, timestamp, caption=None)
```

and the following applies

name	description
urdlist	is the name of the Urd list, and the same urdlist must be specified in both .begin() and .finish(). The urdlist is specified as <user>/<list>, where the <user/> part is optional. The user string is also for authentication, and must correspond to the current URD_AUTH settings, see section B.5.
timestamp	is <i>mandatory</i> , but could be set in either .begin(), .finish(), or both. .finish() will override .begin().
caption	is <i>optional</i> , and can be set in either .begin() or .finish(). .finish() will override .begin().
update	If set to <code>True</code> , the last item in the list may be updated. This option will be discussed in section 7.11.

The Urd transaction database will be written to only when the .finish() function is called. Nothing is stored before calling .finish(), and it is perfectly okay to omit .finish() to avoid storage or during development work.

7.7.1 What if a Build Script is Run Again?

The Urd database is based on transaction log files. There is no way to modify or remove information in the database using the Urd API. The first time a build script is run, it will cause jobs to be built, and if a session is set up using .begin() and .finish, this session will be stored in the Urd database. The second time the same script is run, the Accelerator will look up already built jobs and immediately return job references instead of building anything. As long as there are no changes, the session will look identical to the one stored in the Urd database, and it is possible to execute .finish() without errors. But if there are any discrepancies, such as a job being rebuilt, Urd will refuse to store the differing session and instead cause an exception to notify the user.

Normally, a build script can be written in such a way that re-running it will be consistent with the Urd database and everything is fine. A mismatch with Urd is then an indication of some error. But there are cases when re-writing Urd history is the desired option, and this will be discussed in section 7.11.

7.7.2 Timestamp Definition and Resolution

The “timestamp” used to access items may be stated as either a date, datetime, "datetime", int, (datetime, int), or "datetime+int". Here, "datetime" is a string of format

```
'%Y-%m-%d %H:%M:%S.%f'
```

(See Python’s `datetime` module for explanation.) A specific timestamp could be shorter than the above specification in order to cover wider time ranges. The following examples cover all possible cases.

```
'2016-10-25'           # day resolution
'2016-10-25 15'        # hour resolution
'2016-10-25 15:25'     # minute resolution
'2016-10-25 15:25:00'  # second resolution
'2016-10-25 15:25:00.123456' # microsecond resolution
```

Example of a timestamp with an `int`

```
'2016-10-25+3'
```

Note that

- `ints` sorts *first*,
- `datetimes` without `int` sorts *before* `datetimes` with `ints`,
- shorter `datetime` strings sorts *before* longer `datetime` strings, and
- a timestamp must be > 0 .

7.8 Retrieving an Urd Session

A specific session can be retrieved from the Urd database using its *key* and *timestamp*. There are two sets of functions assigned for this, one that will *record* the lookup to the *ongoing* session, and one that will not. The reason for recording lookups is to provide the possibility for a session to contain references to all other sessions used to create it, and thereby ensuring transparency.

This section is devoted to the function calls that are recorded, and these are `.get()`, `.first()`, and `.latest()`. For any of these calls to work, they have to be issued from *within* a session, i.e. after a `.begin()` call. Otherwise Urd would not be able to record session dependencies. Section 7.9 describes the non-recording versions, and more ways to find sessions and inspect Urd database contents is described in section 7.12.

7.8.1 Finding an Exact or Closest Match: `.get()`

The `.get()` function will return the single session, if available, corresponding to a specified list and timestamp, see the following example

```
urd.begin('ab/anotherlist')
urd.get("ab/test", "2018-01-01T23")
```

The timestamp must match exactly for an item to be returned. If there is no matching item, the `.get()`-call will return an *empty session*, i.e. something like this

```
{'deps': {}, 'joblist': JobList([]), 'caption': '', 'timestamp': '0'}
```

The strict matching behaviour can be relaxed by prefixing the timestamp with one of

“<”, “<=”, “>”, or “>=”.

For example

```
urd.get("ab/test", ">2018-01-01T01")
```

may return an item recorded as 2018-01-01T02. Relaxed comparison is performed “from left to right”, meaning that

```
urd.get("ab/test", ">20")
```

will match the first recorded session in a year starting with “20”, while

```
urd.get("ab/test", "<=2018-05")
```

will match the latest timestamp starting with “2018-05” or less, such as “2018-04-01” or “2018-05-31T23:59:59.999999”.

7.8.2 Finding the Latest Session: `.latest()`

The `.latest()` call will, for a given *key*, return the session with most recent timestamp. For example

```
urd.begin('ab/anothertestlist')
urd.latest('ab/test')
```

will return a complete session, assuming that the list exists and is populated. Otherwise, an *empty session* will be returned.

7.8.3 Finding the first item: `.first()`

The `.first()` function works similarly to `.latest()`, but will instead return the session with the *oldest* timestamp, i.e. the first session stored using the *key*.

7.9 Retrieving an Urd Session without Recording

As pointed out in the previous section, dependencies will be recorded when the `.get()`, `.latest()`, and `.first()` functions are used. If, for some reason, the point is to just have a look at the database to see for example if something is already there and in that case what that is, it can be done using the “peek” functions: `.peek()`, `.peek_first()`, and `.peek_latest()`, like this:

```
urd.peek('test', '2016-10-25')
urd.peek_latest('test')
urd.peek_first('test')
```

There are some good use cases for this, but it is not recommended for general use, they may cause a loss of continuity and visibility.

7.10 Aborting an Urd Session: `.abort()`

There can only be one ongoing Urd session. A new session cannot be started until the current one has finished. There are three ways to end an urd session

- execute the `.finish()` call and have the session recorded. (Or rejected, if a session with the same key and timestamp already exists but with different contents.)
- end the build script “prematurely” before the `.finish()`-call. No data will be stored in Urd.
- issue an `.abort()` call.

The `.abort()` function is used like this

```
urd.begin('test')
urd.abort()
# execution continues here, a new session can be initiated
urd.begin('newtest')
```

Similar to unfinished sessions, aborted sessions will not be stored in the Urd transaction log.

7.11 Truncating and Updating

Since the Urd database is based on transaction log files, it will always keep a consistent history of all events taken place. It is not possible to erase or modify old entries, but it is okay to update the latest item, or set a marker in the log indicating that the list is starting over from a certain date and everything before this marker should not be considered anymore. This makes it possible to both keeping the full history *and* being able to rewrite it. There is full transparency and reproducibility – all sessions before an update or restart marker are always kept in the Urd log file, transaction data is always *appended* to the log files.

7.11.1 Updating the Last Item

To update the last item in a list, set the `update` argument to `True`

```
urd.begin('test', '2016-10-25', update=True)
```

If `update` is `True`, the entry in the test list at '2016-10-25' will be updated, unless the new information is equivalent. The `update=True` call will simply add a new line to the Urd log database, and if the timestamp is the same as the previous entry, the new entry will be selected. Note that this only works for the *last* entry in the database, and that the previous version is not erased, it still exists in the log file.

7.11.2 Truncating a List

Sometimes there is a need to wipe all history or at least start over from an earlier time instance. Since data can not be erased from a transaction log, re-winding to an earlier timestamp is implemented using inserting a “truncation” marker in the log. Any data more recent than what is indicated by the marker is wiped from memory, and sessions with timestamps more recent than the marker can be stored again. The `.truncate()` function is used to insert these markers, and it is used like this

```
urd.truncate(ab/'test', '2016-09-30')
```

This will rollback everything that has happened in the `ab/test` list back to '2016-09-30'. There is also a special case,

```
urd.truncate('ab/test', 0)
```

that will erase all items from memory and cause the list to start over again. Again, remember that Urd is using plain text transaction log that can only be appended to. It is always possible to recover any old result or processing state from these logs.

7.11.3 Truncation Consequences: Ghosts

When a list is truncated, all items after a specified timestamp are made invisible. Assuming that another list has stored a dependency of an item that is truncated, the jobs in this list are now without dependencies that can be looked up. We call them “ghosts”. Ghosts cannot be looked up in Urd, but they are still in the database, where they are marked as ghosts.

7.12 More Search Functions

There are two more functions for finding information in the Urd database: `.list()` and `.since()`.

7.12.1 Listing all Urd Lists: `.list()`

The `.list()` function will return a list of all lists recorded in the database:

```
print(urd.list())
```

may show something like

```
['ab/test', 'ab/live', 'production/import', 'production/live']
```

7.12.2 Listing all Items After a Specific Timestamp: `.since()`

The `.since()` function is used to extract lists of *timestamps* corresponding to recorded sessions. In its most basic form, it is called with a timestamp like this

```
urd.since('ab/test', '2016-10-05')
```

which returns a list with all existing timestamps more recent than the one provided

```
['2016-10-06', '2016-10-07', '2016-10-08', '2016-10-09', '2016-10-09T20']
```

The `.since()` function is rather relaxed with respect to the resolution of the input. The input timestamp may be truncated from the right down to only one digit. An input of zero is also valid. For example, these are all valid

```
urd.since('ab/test', '0')
urd.since('ab/test', '2016')
urd.since('ab/test', '2016-1')
urd.since('ab/test', '2016-10-05')
urd.since('ab/test', '2016-10-05T20')
urd.since('ab/test', '2016-10-05T20:00:00')
```

7.13 Building Jobs: `.build()`

Jobs are dispatched in Urd sessions using the `.build` function. Here is the complete call with all possible parameters.

```
job = urd.build(
    method,
    options={}, datasets={}, jobs={},
    name=None, caption=None,
    why_build=False,
    force_build=False,
    workdir=None,
    concurrency=None,
    **kw
)
```

If `options`, `datasets`, and `jobs` are uniquely defined in the method, they could be entered just as plain keyword arguments. If there are ambiguities, the full `options=` etc. must be used. Here is an explanation of `.build()` parameters:

name	description
method	Name of method to build. Enter <code>'test'</code> here if the method filename is <code>a_test.py</code> . The Accelerator will look for this file in all method directories specified in the Accelerator's configuration file.
kw	Use to specify any unique options, jobs, and datasets without the more elaborate dictionaries described next.
options={}	A <code>dict</code> of options to the method. This overrides options defined in the method itself, but adding options not prototyped in the method is <i>not</i> allowed.
jobs={}	A <code>dict</code> of jobs to the method. It is possible to specify a list of jobs like this: <code>jobs=dict{alljobs=[job1, job2, ...]}</code>

<code>datasets={}</code>	A <code>dict</code> of datasets to the method. Datasets may be lists too, just like <code>jobs</code> above.
<code>name=None</code>	A string associated with the job. Use it to distinguish several jobs created from the same method in <code>joblists</code> .
<code>caption=None</code>	A caption string. For decorative purposes only, this has no practical use.
<code>why_build=False</code>	Don't build this job. If it appears to need building stop and print information on how this build could have been avoided. This is a debugging aid.
<code>force_build=False</code>	Build this job even if there appears to be no reason to. This is also mostly a debugging aid.
<code>workdir=None</code>	If specified, the job will be built in this <code>workdir</code> , assuming the <code>workdir</code> is specified in the configuration file as <i>either</i> source or target.
<code>concurrency=None</code>	Specify max number of parallel processes. This can be used to save memory, see section A.5 for more information.

The `.build()` function will only build a job when it has to, otherwise it will just return a job reference to an existing matching job. In order to match, an existing job must have

- exactly the same source code, i.e. the *hash* of the source code must match,
- exactly the same options, datasets, and jobs.

If the source code is changed, a job rebuild can be prevented using the `equivalent_hashes` variable as explained in section 4.4.3.

7.13.1 Building Chained Jobs: `urd.build_chained()`

This is a special version of `.build()` that can be used for linking a set of dataset-creating jobs. This function was created for the purpose of having build scripts that imported a large set of files in a `for`-loop. It is used like this

```
def main(urd):
    # Import a list of files
    for filename, timestamp in listoffiles:
        urd.begin('import', timestamp)
        urd.latest('import')
        job = urd.build_chained(
            'csvimport',
            filename=filename,
            ...
            name='importing',
        )
        job = urd.build_chained(
            'dataset_type',
            source=job,
            ...
            name='typing',
        )
        urd.finish('import')
```

This example will build a chain of `csvimport` jobs, and one chain of `dataset_type` jobs. Each `dataset_type` job will have a corresponding `csvimport` dataset as source. The `.build_chained()` function works, provided that

- the method to build has a dataset named `previous`,
- a unique `name=` is set in `.build_chained()`, and

- `urd.latest()` is called inside the Urd session.

The call to `urd.latest()` is necessary for the dependency-logic to work, but the output from the call can be discarded.

7.14 Changing workdir: `.set_workdir()`

The target workdir specified in the configuration file is the only workdir that is written to by default. Any other workdir is read only. This behaviour can be overridden, either

per job, using the `workdir=...` option to `urd.build()` as shown in section 7.13, or using `urd.set_workdir()`.

The latter,

```
def main(urd):
    urd.set_workdir(<workdir>)}
```

will set the workdir for all coming `.build()` calls in the current build script. It can still be overridden using the `workdir=` option to `urd.build()`.

7.15 Profiling a Build Script: `.print_exeetimes()`

The `JobList` object has a helper function that can be used to print profiling information for the joblist. The following example is self-explanatory

```
def main(urd):
    ...
    urd.joblist.print_exeetimes()
```

This will print execution times for all jobs in the session to `stdout`. It may for example look like this

```
Time per method:
  color2          23.7 seconds (25%)
  csvexport       17.5 seconds (18%)
  lowpass2        15.6 seconds (17%)
  newcol          14.4 seconds (15%)
  black           5.7 seconds (6%)
  colimage        5.4 seconds (6%)
  sync            4.7 seconds (5%)
  clamp           3.8 seconds (4%)
  dataset_type    2.7 seconds (3%)
  csvimport       1.4 seconds (1%)
Total time 94.8 seconds
```

The methods are sorted by execution time, top to bottom.

7.16 Passing Flags from the Command Line

Flags can be added to a build script at run time using a comma separated list like this

```
ax run [script] --flags=verbose,skiptest
```

The flags will appear in the `urd`-object like this

```
def main(urd):
    if 'verbose' in urd.flags:
        print('verbosity')
```

7.17 The `ax urd` Command and the Board

In later versions of the Accelerator, the `ax urd` command provides a simple way to visualise what is stored in the Urd database.

In addition, the Accelerator *board* server will render a link “urd” in the top right corner of the web page that can be used to explore the Urd database in a web browser.

7.18 Multi-User Urd Consistency

Urd can be accessed by a large number of clients. Each client may add to or truncate any list at any time. In order to avoid race conditions and make the database deterministic, all `add-` and `truncate-`requests appears in a sequential manner to the Urd server. Each request is assigned a unique timestamp, and stored in the requested list.

When Urd is restarted, it reads all the database files, and sorts all rows in order of the receive timestamp. Thereafter, each row is applied in increasing time order to the internal, RAM-based database. Due to the unique timestamping, the result is a deterministic replica of the previous run.

Chapter 8

Standard Methods

The Accelerator is shipped with a set of common standard methods, including methods to import, type, export, and hash partition data. These methods are found in the method directory `./standard_methods`. All methods in `standard_methods` are designed and tested to work on both Python2 and Python3.

8.1 csvimport – Importing Data Files

The `csvimport` method is used to import a text files into a dataset. Data is imported one row per slice in a round robin fashion. Input data is assumed to be in a tabular format, i.e. it is composed of a number of rows, each having the same number of columns separated by a separator token. A common format of this type is the Comma Separated Values (CSV) format, but `csvimport` is much more flexible, as seen in the table of options below. For example, `csvimport` can handle any separator character, skip or parse labels on the first line, and supports advanced quote support. It also deals with “broken” input data in a predicted and user controllable way. `csvimport` jobs can be chained, so any number of input files can be connected in a dataset chain.

8.1.1 Options

name	default	description
<code>filename</code>	<i>mandatory</i>	Name of file to import. The filename is mandatory and the file may either be a <i>plain text</i> file or a <i>gzipped</i> file. It is also possible to specify a filename including a path. If the path begins with a slash, it is absolute. Otherwise, the path is relative to the “ <code>input directory</code> ” configuration parameter specified in the Accelerator’s configuration file, see section B.4. A relative path makes it possible to relocate files to a different directory without triggering job remake.
<code>separator</code>	<code>,</code>	Field separator character. Accepts a single <code>iso-8859-1</code> character. Leave this empty to import each line of the input file into a single column.
<code>comment</code>	<code>”</code>	Lines beginning with this character are ignored. Accepts a single <code>iso-8859-1</code> character, or the empty string for no comments. Commented lines are stored in the <code>skipped</code> dataset.
<code>newline</code>	<code>”</code>	Newline character. Empty means <code>”\n”</code> or <code>”\r\n”</code> . Alternatively any single <code>iso-8859-1</code> character can be chosen.
<code>quotes</code>	<code>”</code>	Quote character. Empty or <code>False</code> means no quotes, <code>True</code> means both <code>’</code> , and <code>”</code> , any other character means itself. This supports standard CSV escaping, i.e. doubling the quote character. <code>’a’b’</code> → <code>a’b</code> Unquoting applies to any field that starts with a quote character, quote characters inside other fields are not treated specially.
<code>labelsonfirstline</code>	<code>True</code>	If set to <code>True</code> , data on the first line of the file will be used as column labels. These will be decoded with <code>str.decode(’utf-8’, ’backslashreplace’)</code> . If <code>False</code> , labels must be entered using the <code>label</code> option, see <code>labels</code> below.
<code>labels</code>	<code>[]</code>	If <code>labelsonfirstline</code> (see above) is set to <code>False</code> , labels must be provided using this option. For example <code>labels = [’foo’, ’bar’,]</code> . If specified, these labels always apply, regardless of <code>labelsonfirstline</code> .
<code>rename</code>	<code>{}</code>	This option makes it possible to change the column names read from the first line of the input file. Renaming happens before discard. It accepts a dictionary <code>{old_name: new_name,}</code> .

<code>lineno_label</code>	<code>''</code>	If set, <code>lineno_label</code> becomes a column containing line numbers. Line numbers start at one (1), and corresponds to line numbers in the input file.
<code>discard</code>	<code>set()</code>	Labels in the discard set will not be stored in the dataset.
<code>allow_bad</code>	<code>False</code>	By default, this is set to <code>False</code> and an error will be asserted if there are problems parsing the input data, see section 8.1.3. Setting it to <code>True</code> will put all “bad” lines together with the corresponding line numbers into a separate dataset named <code>bad</code> . It is recommended to check the resulting datasets if enabling this option!
<code>skip_lines</code>	<code>0</code>	Skip this many lines at the start of the file. This is useful for data files that starts with a header, for example. Skipped lines will be stored in the <code>skipped</code> dataset.
<code>allow_extra_empty</code>	<code>False</code>	Still consider a line good if it has extra empty fields at the end.
<code>skip_empty_lines</code>	<code>False</code>	Ignore empty lines.
<code>strip_labels</code>	<code>False</code>	Do <code>.strip()</code> on all labels (happens before rename).

8.1.2 Datasets

name	default	description
<code>previous</code>	<code>None</code>	Previous dataset if creating a chain.

8.1.3 Bad Lines

A line is flagged as “bad” for one of two reasons

- there is a problem with quoting, or
- there is an incorrect number of separators.

for example

```
"a","b" "c"      # invalid assuming two or three comma separated columns
"a","b"" "c"     # valid assuming two comma separated columns
```

8.1.4 Output

The result of the `csvimport` is a dataset named `default`. Lines marked as “bad” will be stored in the dataset `bad`, while skipped and commented lines will be stored in the dataset `skipped`. All columns (except the `lineno` column) will be of type `bytes`. Typically, the dataset from `csvimport` is fed to a `dataset_type` job for column typing.

8.1.5 Line Numbers

A column with line numbers is always attached to the `bad` and `skipped` datasets, and conditionally using `lineno_label` to the main dataset. Line numbers start at one (1), and always corresponds to the lines in the input dataset. For example, if there are labels on the first line of the input file, this line is number 1. Any line number can thus only appear in one of the main, `bad`, or `skipped` datasets.

8.1.6 Limitations

Each line is limited to slightly less than 16MB. However, this is just a constant in the code (`BIG_Z` in `standard_methods/csvimport.py`) that is by default set to a value that allows the Accelerator to run on low memory platforms (it will allocate $sliceno * 2 + 3$ buffers of this size). If you have enough memory you can set it as high as $2^{30} - 1 = 1073741823$ (1GB), but higher than that will not work correctly.

8.1.7 Example Invocation

An example invocation is the following

```
urd.build(csvimport', filename='inputfile.txt', separator='\0')
```

this will import the file `inputfile.txt` assuming that there are labels on the first line and the column separator is a null character (`0x00`, `'\0'`).

8.2 csvimport_zip – Importing zip Archives

The `csvimport_zip` method is a wrapper around `csvimport` that is used to import files stored in zip archives. One or more files in a zip archive can be imported by a call to this function, and each file will be imported to a separate dataset.

8.2.1 Options

All options to `csvimport` are available to this method as well, and the `filename` option is used to specify the name of the zip file.

name	default	description
<code>inside_filenames</code>	<code>{}</code>	Dictionary from filename in zipfile to dataset name, <code>{'filename in zip': 'dataset name', ...}</code> . If left empty, all files will be imported to datasets with cleaned up names. If there is only one file imported from the zip (whether specified explicitly or because the zip only contains one file) this will also end up as the default dataset.
<code>chaining</code>	<code>on</code>	Can be one of: <code>off</code> – Do not chain the imports. Use <code>datasets.previous</code> on all datasets. <code>on</code> – This is the default, and means <code>by_dict</code> if <code>inside_filenames</code> is set and <code>by_ziporder</code> if it is not. <code>by_filename</code> – Chain the imports in filename order. <code>by_dsname</code> – Chain the imports in dataset name order. <code>by_dict</code> – Chain the imports in the order in <code>inside_filenames</code> . Note that you must pass <code>inside_filenames</code> an <code>OrderedDict</code> for this. <code>by_ziporder</code> – Chain the imports in the order the files are in the zip file.
<code>include_re</code>	<code>''</code>	Regex of files to include, matches anywhere.
<code>exclude_re</code>	<code>''</code>	Regex of files to exclude, takes priority over <code>include_re</code> .
<code>strip_dirs</code>	<code>False</code>	Strip directories from filename (<code>a/b/c</code> → <code>c</code> .)

If chaining is enabled, the last dataset will be the `default` dataset. Note that naming a non-last dataset “`default`” is an error. If `strip_dirs` is set, the filename (as used for both sorting and naming datasets, but not when matching regexes) will not include directories. The default is to include directories.

8.2.2 Example invocation

```
urd.build(  
    "csvimport_zip",  
    filename="data_Q2_2019.zip",  
    exclude_re=r"(__MACOSX|\.DS_Store)",  
    chaining="by_filename",  
    strip_dirs=True,
```

DRAFT

8.3 dataset_type – Typing Datasets

The `dataset_type` method will read a source dataset or dataset chain and type its columns. This method is primarily used for typing datasets created by `csvimport`, but it can type any column of type `bytes`, `ascii`, or `unicode` to any other type.

The method will also hash partition the output dataset if the `hashlabel` input parameter is set, causing a new dataset to be created. For additional information about hash partitioning, see the `dataset_hashpart` method in section 8.5.

The default behaviour is to append new columns with typed data to the existing source dataset. These columns will have the same name as the untyped version of the data, making the untyped data “inaccessible”, even if it is still in the dataset. Using the `rename` option, typed columns can be assigned a name that differs from the original name, so that both typed and untyped data are available simultaneously. This brings transparency to the typing process. (However, even if the untyped data is “inaccessible” in the typed dataset, it is still available if referenced as the input dataset.)

In order to type the data, the input data is subject to parsing. Some datasets may contain data that is incorrect in the sense that it causes parsing errors when typing. Unparseable data can either be replaced by a default value or removed from the dataset. Since the Accelerator’s dataset type does not permit removal of rows, i.e. datasets can not shrink, `dataset_type` will in this situation create a new dataset containing only the rows containing typeable data.

If typing a dataset chain, any columns that do not have the same type over all the typed datasets will be discarded.

8.3.1 Datasets

name	default	description
source	<i>mandatory</i>	Dataset to type.
previous	None	Previous dataset if creating a chain.

8.3.2 Options

name	default	description
column2type	<code>{}</code>	A dictionary from column label to type, for example <code>{'movie': 'unicode:UTF-8',}</code> .
timezone	None	Set input timezone for <code>datetime</code> columns. Output will be in UTC. See text.
hashlabel	None	Hash partition dataset based on this column. Leave as None to inherit hashlabel, set to <code>''</code> to not have a hashlabel. Hashing causes a new dataset to be created.
defaults	<code>{}</code>	A dict from column name to default value, for example <code>{'COLNAME': value}</code> . Method will fail if data is unconvertible unless <code>filter_bad = True</code> .
rename	<code>{}</code>	A dictionary from old name to new name, for example <code>{'old': 'new'}</code> The old name and data will be preserved, unless a new dataset is created, and the column with the new name will contain the typed data.
caption		Optional caption. A reasonable caption is created automatically if left blank
discard_untyped	None	If set to True , force creation of new dataset and make untyped columns inaccessible. If set to False , an error is generated if any columns were not preservable.

<code>filter_bad</code>	<code>False</code>	If <code>False</code> , fail when a value fails to convert and there is no default. If <code>True</code> , filter out the line with the unconvertable value. This will create a new dataset.
<code>numeric_comma</code>	<code>False</code>	If <code>True</code> , write decimal number as “3,14” instead of default “3.14”.
<code>length</code>	<code>-1</code>	Go back at most this many datasets. The default is <code>-1</code> , which goes until <code>previous.source</code> if it exists, or first dataset in chain otherwise.
<code>as_chain</code>	<code>False</code>	If hash partitioning, avoid re-writing at the end by doing one dataset per slice.
<code>compression</code>	<code>6</code>	<code>gzip</code> compression level.

The `timezone` option is used to specify which input timezone a `datetime` column is in. It does not work for `date` or `time` columns. It can be set to anything accepted by the system’s `$TZ`. Note

- It does not work for `%s` (which is always in UTC).
- No error checking can be done on this (`tzset(3)` can not return failure).
- On most 32bit systems this will break dates outside about 1970 - 2037.
- Setting this will mask most bad dates (`mktime(3)` “fixes” them).
- Do not set this to `'UTC'`, leaving it as `None` is faster and safer.

8.3.3 Example Invocation

An example invocation is the following

```
urd.build('dataset_type',
    source=...,
    previous=...,
    column2type=dict(
        auct_start_dt='datetime:%Y-%m-%d',
        brand='json',
        item_id='number',
        comp='unicode:utf-8',
    ),
)
```

8.3.4 Typing

This section describes all typing possibilities in detail. Default behaviour when typing numbers (i.e. `floats`, `ints`, and `numbers`) is that any number of whitespaces before and after the actual number are silently discarded.

Numbers

The `number` type is integer or floating point.

<code>number</code>	<code>int</code> or <code>float</code>
<code>number:int</code>	<code>int</code> , but accepts float-formatted ints.

Floats have standard double precision (64 bits), ints can be up to $\pm(2^{1007} - 1)$.

Integers are enforced using `number:int`, and the type accepts trailing decimal zeroes like 7.0, 4.000 etc. This is useful when typing datafiles where numbers actually are integers but have trailing zero decimals. It is especially useful if the numbers are too big to be represented exactly as doubles (i.e. more than 2^{53}).

Floating Point Numbers

Floating point numbers may be stored as 32 or 64 bits. The *ignore* option ignores any trailing characters after the number.

<code>float32</code>	<code>float64</code>	<i>default</i>
<code>float32i</code>	<code>float64i</code>	<i>ignore</i> , will discard trailing garbage

Integers

Integers are stored as either 32 or 64 bits. Parsing takes base into account, so in addition to decimal numbers, it is also straightforward to parse octal and hexadecimal numbers. The *ignore* option causes parsing to ignore trailing garbage characters.

<code>int32_0</code>	<code>int64_0</code>	<i>automatic base detection</i> , consider avoiding this in favor of a deterministic type
<code>int32_0i</code>	<code>int64_0i</code>	<i>automatic base detection</i> , ignore trailing garbage
<code>int32_8</code>	<code>int64_8</code>	<i>octal</i>
<code>int32_8i</code>	<code>int64_8i</code>	<i>octal</i> , ignore trailing garbage
<code>int32_10</code>	<code>int64_10</code>	<i>decimal</i>
<code>int32_10i</code>	<code>int64_10i</code>	<i>decimal</i> , ignore trailing garbage
<code>int32_16</code>	<code>int64_16</code>	<i>hexadecimal</i>
<code>int32_16i</code>	<code>int64_16i</code>	<i>hexadecimal</i> , ignore trailing garbage

Convert to Boolean

It is common that a column holds values that are to be interpreted as either `False` or `True`. The following types handles strings and floats.

<code>strbool</code>	<code>False</code> if value in (<code>False</code> , <code>0</code> , <code>f</code> , <code>no</code> , <code>off</code> , <code>nil</code> , <code>null</code> , <code>""</code>) <code>True</code> otherwise
<code>floatbool</code>	<code>True</code> when float has bits set. Is <code>False</code> otherwise.
<code>floatbooli</code>	same + <i>ignore</i> trailing garbage

Time and Date

There are three types relating to time available, `date`, `time`, and `datetime`. Each of these has a corresponding version that ignores trailing garbage characters. All time types require a format specification (represented here by an asterisk) as described below

<code>date:*</code>	a date with format specifier
<code>datei:*</code>	same + <i>ignore</i> trailing garbage
<code>time:*</code>	a time with format specifier
<code>timei:*</code>	same + <i>ignore</i> trailing garbage
<code>datetime:*</code>	a date + time with format specifier
<code>datetimei:*</code>	same + <i>ignore</i> trailing garbage

The format is standard Python time formats, like shown in these examples

```
# will match for example '2017-03-22'  
auct_start_dt='date:%Y-%m-%d'  
# will match for example '183000', i.e. half past six in the evening  
tod='time:%H%M%S'
```

```
# will match for example '2017-03-22 18:30:15'  
timestamp='datetime: '%Y-%m-%d %H:%M:%S'
```

Strings and Byte Sequences

There are a number of ways to read string and byte data, depending on how the raw input data is to be interpreted. The basic types are shown first, and the more advanced variations and options will be described below.

<code>bytes</code>	list of bytes
<code>bytesstrip</code>	list of bytes, strip ASCII characters 8-13, 32 from start and end
<code>ascii</code>	list of ASCII characters
<code>asciistrip</code>	list of ASCII characters, strip characters 8-13, 32 from start and end

When typing to unicode and ASCII, there are several ways to handle individual unparseable characters. For unicode, there are two types,

<code>unicode:*</code>	list of unicode characters
<code>unicodestrip:*</code>	list of unicode characters, strip characters 8-13, 32 from start and end

The asterisk represents options that take the form

```
"codec" #or  
"codec/errors"
```

`unicode:codec/errors` will read bytes encoded in `codec` and write "unicode" (which is stored as utf-8, but that's invisible to the Python side). `codec` is often `utf-8`, but could be for example `utf-8`, `ascii`, `iso-8859-1`, `iso-8859-15`, `cp437`, or `windows-1252` etc. See the Python documentation

<https://docs.python.org/3/library/codecs.html#standard-encodings>

for more information. The `errors` part is optional, and can be one of

<code>strict</code>	The default, an error marks this row as bad
<code>ignore</code>	All unparseable bytes are discarded.
<code>replace</code>	All unparseable bytes are replaced by the unicode replacement character (" <code>\ufffd</code> ").

Using `strict` will cause errors if unparseable. For example, typing the string "`ab\xffc`" will give an error (`strict`), "`abc`" (`ignore`), or "`ab\xffdc`" (`replace`). `strip` will happen before `ignore`.

ASCII is similar, there are two types

<code>ascii:*</code>	list of ASCII characters
<code>asciistrip:*</code>	list of ASCII characters, strip characters 8-13,32 from start and end

where the argument is one of

<code>strict</code>	The default, an error marks this row as bad
<code>ignore</code>	All unparseable bytes are discarded
<code>replace</code>	All unparseable bytes are replaced by an octal escapes " <code>\ooo</code> "
<code>encode</code>	Like <code>replace</code> except " <code>\</code> " is also replaced by " <code>\134</code> " (for full reversibility).

Using `strict` will cause errors if unparseable. `strip` will happen before `ignore`.

8.4 csvexport – Exporting Text Files

The `dataset_export` method is used to export datasets to column based text files (CSV, Comma Separated Values). It can export plain files and `gzip`-compressed files, export a chain of datasets, export one output file per slice, and more. Read the Options section for full details.

Options

name	default	description
<code>filename</code>	<code>result.csv</code>	Name of output file. File will by default be stored in the job's job directory. The filename has to end with <code>".csv"</code> for plain text files, and <code>".gz"</code> for gzipped output.
<code>separator</code>	<code>','</code>	Column separator.
<code>labelsonfirstline</code>	<code>True</code>	If <code>True</code> , write column names on first row.
<code>labels</code>	<code>[]</code>	Specify which labels to export. An empty list corresponds to all labels in dataset.
<code>chain_source</code>	<code>False</code>	If <code>True</code> , read a dataset chain from <code>datasets.source</code> back to <code>jobs.previous</code>
<code>quote_fields</code>	<code>empty string</code>	Export quoted fields. Must be empty (no quote character, default), <code>"'"</code> , or <code>""</code> .
<code>lazy_quotes</code>	<code>True</code>	Only quote field if value needs quoting.
<code>line_separator</code>	<code>'\n'</code>	Line separator.
<code>none_as</code>	<code>None</code>	A string or <code>{label: string}</code> to use for <code>None</code> -values. Default <code>'None'</code> (<code>'null'</code> for json).
<code>sliced</code>	<code>False</code>	Each slice is exported in a separate file when <code>True</code> . If so, use <code>"%02d"</code> or similar in filename as placeholder for the slice number.
<code>compression</code>	<code>6</code>	gzip level

Datasets

name	default	description
<code>[source,]</code>	<code>mandatory</code>	Either A <i>single</i> dataset or a <i>list</i> of datasets.

Jobs

name	default	description
<code>previous</code>	<code>None</code>	Job reference to previous <code>csvexport</code> if chained.

8.4.1 Example Invocation

An example invocation is the following

```
urd.build(  
    'csvexport',  
    source='test-3/foo',  
    filename='output.txt.gz',  
    separator=' ',
```

```
) quote_fields="'",
```

DRAFT

8.5 dataset_hashpart – Hash Partition a Dataset

The `dataset_hashpart` method will create a new dataset based on its `source` dataset. The new dataset will be hash partitioned on a column specified in the options.

Options

name	default	description
<code>hashlabel</code>	<i>mandatory</i>	column for hashing, required. Note that columns typed as <code>list</code> , <code>set</code> , or <code>json</code> cannot be used for hashing.
<code>length</code>	<code>-1</code>	Go back at most this many datasets in a chain. Default is <code>-1</code> , which goes back to <code>previous.source</code> if it exists, or to the first dataset in the chain otherwise.
<code>caption</code>		Optional caption. A reasonable caption is created automatically if left blank
<code>as_chain</code>	<code>False</code>	True generates one dataset per slice, False generates one dataset. Default <code>False</code> .

Datasets

name	default	description
<code>source</code>	<i>mandatory</i>	Source dataset to hash partition
<code>previous</code>	<code>None</code>	Previous dataset to chain to.

8.5.1 Example Invocation

An example invocation is the following

```
urd.build('dataset_hashpart', source=jid, hashlabel='start_date')
```

8.5.2 Hashing Details

This method will create a new dataset based on all the data in the source dataset. The difference between input and output is in which slices the rows will be stored. For each row, the target slice is determined based on the output value of a hashing function applied to a certain column (the `hashlabel`) of that row. To illustrate the operation, the code is similar to

```
from accelerator.gzutil import siphash24

target_sliceno = siphash24(cols[hashlabel]) % params.slices
```

The `hashlabel` column can be of any type *except* `json` and `pickle`!

8.5.3 Notes on Chains

1. The default operation is to hash partition a complete chain of datasets from `source` back to `previous.source`. This is controlled by the `length` option.
2. Internally, `dataset_hashpart` always generates one dataset per slice in a chain. This is also what is returned if `as_chain == True`. Otherwise, all datasets will be concatenated into one. Thus, there is a choice of either having the output as a chain of datasets – or as a single dataset. The chain will execute faster, since the concatenation step is omitted.

8.6 dataset_filter_columns – Removing Columns from a Dataset

The `dataset_filter_columns` method removes columns from a dataset. It is typically run before applying methods that operate on all columns of a dataset and only a subset of the columns are required. A typical example is `dataset_hashpart` that operates on all columns of a dataset. If not all columns are needed, time and storage can be saved by removing columns using this method prior to applying `dataset_hashpart`.

Note that this method only updates soft links, and no data is actually copied. So execution time is typically a fraction of a second and no redundant data is written to disk.

Options

name	default	description
<code>keep_columns</code>	<code>[]</code>	A list of columns to keep.
<code>discard_columns</code>	<code>[]</code>	A list of columns to remove.

Only one of `keep_columns` and `discard_columns` may be populated.

Datasets

name	default	description
<code>source</code>	<i>mandatory</i>	Source dataset to filter.
<code>previous</code>	<i>None</i>	Previous dataset to chain to.

8.7 dataset_sort – Sorting a Dataset

The method `dataset_sort` is used to sort relatively large datasets. One or more columns may be selected for sorting, and it will sort one column at a time. The sorting algorithm is stable, meaning that things with equal sorting keys will keep their order.

Options

name	default	description
<code>sort_columns</code>	<i>mandatory</i>	A column or a list of columns. If a list is specified, sorting will be carried out from left to right.
<code>sort_order</code>	<code>ascending</code>	Could be reversed by specifying <code>descending</code>
<code>sort_across_slices</code>	<code>False</code>	If <code>False</code> , only sort within slices. Otherwise sort across slices.
<code>trigger_column</code>	<code>''</code>	If <code>sort_across_slices</code> is set, use this to delay the slice switches to the next line where the value in this column changes.

Datasets

name	default	description
<code>source</code>	<i>mandatory</i>	A dataset to sort.
<code>previous</code>	<code>None</code>	A previous dataset to chain to.

8.7.1 Sorting `None` and NaN values

The special values `None` and NaN follow these rules

- NaN will sort same as `+Inf`, i.e. *last*.
- `None` sorts as `-Inf`, i.e. *first* in float columns. Intermingled `None` and `-Inf` will keep their original order due to the stable sorting algorithm.
- `None` sorts *last* in date, time, and datetime columns.
- For all other types, `None` sorts *first*.

8.7.2 Spreading of Left Over Values

If the number of rows in a dataset is not even divisible by the number of slices, some slices will have one more row than others. Instead of putting this data in, say, the first slices, `dataset_sort` attempts to even out any bias by selecting the slices that get the additional data row in a pseudo-random manner. In order to have the sorting stable, selection of slices is based on the first values of the sorting column. It is not perfect, if the data is already sorted the first slices will be picked, but it is stable, which is the most important thing.

8.7.3 The Trigger Column

If sorting across slices, i.e. the full dataset will be sorted from the first to the last slice, the `trigger_column` could be used to delay the slice switching so that it does not appear unless the value in the column is changing. This is beneficial in parallel processing tasks, because it localises values to single slices.

8.7.4 A Practical Limitation

`dataset_sort` may run out of memory, depending on the amount of RAM, number of parallel processes, and the input dataset. If this should happen, one simple solution is to use the `CONCURRENCY` option to reduce the number of concurrent parallel processes, for the `dataset_sort` method only. See section A.5.

Discussion: internally, the method works by reading the columns to sort by, and create a sorting indexing column that stipulates the sorting order. Each column is then read, one at a time, and written according to the sorting column. All columns except the sorting columns are read as raw data, meaning that no de-serialising into Python's internal format is required. This makes sorting very fast and relatively low on memory useage. However, since it has to hold the index as well as one full column of data in memory, the method has in theory limited sorting capability. Still, a standard computer can sort very large datasets without trouble.

DRAFT

8.8 dataset_checksum, dataset_checksum_chain

The `dataset_checksum` method is used to create a single checksum from a dataset based on one or more columns. The chained version returns a single checksum from a dataset chain. It is mainly intended as a debugging aid, enabling comparison of datasets across machines, even if they have different slicing.

If `options.sort=False`, hashing will depend on the actual row order of the dataset. If, on the other hand, `options.sort=True`, hashing will be *slice invariant* and *row order invariant*, meaning that the methods only look at the contents of the dataset(s).

Chain limits will affect the checksum of a chain, so if checksumming two chains containing the same data, but with different number of chained datasets, their checksums will differ.

Note that sorting uses about 64 bytes per row, upper limiting the size of hashable datasets. This corresponds to about 1GB of RAM per 20 million lines or so.

Note also that `pickle`-columns cannot be used for checksumming, since the order of the pickled data structure is not defined.

The checksum will be printed to `stdout` as well as written to `result.pickle`.

Options

name	default	description
<code>columns</code>	<code>set()</code>	A set of columns to base the checksum on. Leave blank for all columns
<code>sort</code>	<code>True</code>	Sort dataset before hashing, see text.
<code>chain_length</code>	<code>-1</code>	Number of datasets in chain to hash. Only for <code>dataset_checksum_chain</code> .

Datasets

name	default	description
<code>source</code>	<i>mandatory</i>	A dataset to checksum.
<code>stop</code>	<code>None</code>	Stop hashing at this dataset. Only for <code>dataset_checksum_chain</code> .

8.9 dataset_merge – Merge Several Datasets into One

Merge two or more datasets. The datasets must have the same number of lines in each slice and if they do not have a common ancestor you must set `allow_unrelated=True`. Columns from later datasets override columns of the same name from earlier datasets. Note that merging is a very fast constant time operation.

Options

name	default	description
<code>allow_unrelated</code>	<code>False</code>	Must be <code>True</code> to join datasets that do not share a common ancestor.

Datasets

name	default	description
<code>[source,]</code>	<i>mandatory</i>	A list of datasets to merge.
<code>previous</code>	<code>None</code>	Previous for the merged dataset.

8.10 dataset_unroundrobin – Transpose a Dataset

This method will transpose the dataset so that first row from slices 0, 1, 2, ... will end up in slice 0 rows 0, 1, 2 ... See figure 5.2 on page 60. The method always creates a new dataset. If `trigger_column` is set to the name of a column, slice switching is delayed until the data in the specified column changes. Otherwise, the new dataset will have the same number of rows per slice as the source dataset.

The `trigger_column` will ensure that no value is spread over more than one slice, assuming that all occurrences of the same value are grouped together, which is helpful in many parallel processing cases.

Options

name	default	description
<code>trigger_column</code>	<code>None</code>	Set to delay slice switching based on content.

Datasets

name	default	description
<code>source</code>	<i>mandatory</i>	An input dataset.
<code>previous</code>	<code>None</code>	Previous for the created dataset.

Chapter 9

Accelerator Command Line Tools

DRAFT

The Accelerator is controlled using the `ax` shell command. In order to run any command, `ax` needs to have access to a configuration file (see section B.4). The `ax` command will look for this file first in the current directory, and then recursively in directories above it.

It is assumed that the Accelerator server and build script `run` commands are executed from the same directory. This will work out of the box. But if set up correctly, they could be run from different directories or even from different computers if necessary.

Asking for help is always an option. To begin,

```
ax --help
```

will print something like

```
usage: ax [--config CONFIG_FILE] command [args]

optional arguments:
  -h, --help            show this help message and exit
  --config CONFIG_FILE  configuration file
  --version              alias for the version command

commands:

  abort  abort running job(s)
board-server  runs a webserver for displaying results
  curl  http request (with curl) to urd or the server
  ds    display information about datasets
  grep  search for a pattern in one or more datasets
  init  create a project directory
  intro show introduction text
  job   information about a job
method  information about methods
  run   run a build script
script  information about build scripts
server  run the main server
  urd   inspect urd contents
urd-server  run the urd server
  version  show installed accelerator version
workdir  information about workdirs

aliases:
  cat = grep -e ""

use "ax <command> --help" for <command> usage
try "ax intro" for an introduction
```

All commands accept the option `--help` to display more detailed information about its usage.

All boolean arguments (e.g. `--headers`, `--allow-missing-columns`) have corresponding `--no-/--not-/--dont-` versions (e.g. `--no-headers`, `--dont-allow-missing-columns`).

Next, a quick summary will be presented before going into the details of each command.

9.1 Quick Summary of Commands

To start with, a new project is initiated using the `init` command,

```
ax init
```

This command creates everything required for starting a project, including directories for input data and `workdir` storage. It even sets up an empty `git` repository for the project's source files. All this can be overridden or specified in detail using the command's switches.

Next, an Accelerator server is started using

```
ax server
```

The server is configured using the configuration file that is typically set up by the `init` command.

To see a list of available build scripts and their description, type

```
ax script
```

A build script is run using

```
ax run
```

The idea is that the *main* project build script is run just like that, but it is possible to specify specific scripts to run as arguments. To run a script named `build_name.py`, type `ax run name`. For systems with limited memory, an interesting idea is the possibility to set the *concurrency* for all or some methods using a switch. This will save memory resources, since only a subset of all parallel processes will be run at once.

At any point in time, it is possible to abort the server processing using

```
ax abort
```

When coding, it is useful to have access to the list of available methods and their interfaces, this is provided using the

```
ax method
```

command. It can either provide a list of methods per package, together with an optional description, or given a specific method name it will print its documentation + interface.

When a script has been run, any job can be inspected using

```
ax job <jobid/job-spec>
```

that will print execution times, input parameters, output file names etc for jobs. See section ?? for alternative ways to specify a job.

Similarly,

```
ax ds <dataset/dataset-spec>
```

holds the key for a lot of information about datasets. See section ?? for alternative ways to specify a dataset. Everything from column names, max/min-values, and types to distribution of elements per slice is available. If the input is a job, it will look for the job's default dataset, but it can also show the name of all dataset in a job, and more.

To look at the data in a dataset, use

```
ax grep <pattern> <dataset>
```

or just

```
ax cat <dataset>
```

to stream the data to `stdout`. Using switches data can be output per slice or “in parallel”, it can print header labels and more.

To learn about available workdirs, and which jobs that exists in workdirs,

```
ax workdir
```

is used to either list all available workdirs, or list all jobs in a workdir.

Finally,

```
ax version
```

will print the current version of the Accelerator.

9.2 Job and Dataset References for Command Line Tools

For commands such as `ax job` and `ax ds`, there are several convenient ways to specify the input jobs or datasets.

by **jobid**, like this

```
ax job dev-12
```

by **method name**, like this

```
ax job csvimport
```

This will look up the *latest* and *current* job created by the `csvimport` method, i.e. the latest finished job with current (unmodified) source code.

by **urdlst** key/timestamp, for example like this

```
ax job :import/2021-09-09:
```

Here are examples of to use this:

```
:import: latest job in joblist with latest timestamp in import.
:import/2021-09-09: latest job in joblist with given timestamp.
:import/2021-09-09:csvimport latest csvimport job in joblist with timestamp.
:import/2021-09-09:-2 second to last job in joblist with timestamp.
```

In addition, the `~` and `^` can be used to “follow” the history of jobs in a similar fashion as is provided by the `git` version control tool. Here are some examples of how to use it

```
:import/2021-09-09:~~ corresponds to the latest job in import that is two
timestamps back.
:import/2021-09-09:~2 is the same thing
:import/2021-09-09:^ looks up the .previous to latest job in joblist with
timestamp.
```

Note that for specifying *datasets*, the `~` and `^` must be suffixed to a dataset specifier, and not a job specifier. In practice, this means that a dataset name has to be given explicitly, i.e. write `job-42/default~` and *not* `job-42~`.

9.3 Initialisation: `ax init`

In order to start a new project, a few things need to be set up, in particular

- identify existing / create new *workdirs*,
- identify existing / create new *method packages*, and
- write a *configuration file*.

This can be done manually, but a simple way to start from scratch is to use the `init` command

```
ax init
```

with the following options

```
ax init --help
```

```
usage: ax init [-h] [--slices SLICES] [--name NAME] [--input INPUT] [--force]
              [--tcp [HOST/PORT]] [--no-git] [--examples] [DIR]
```

```
creates an accelerator project directory.
defaults to the current directory.
creates accelerator.conf, a method dir, a workdir and result dir.
both the method directory and workdir will be named <NAME>,
"dev" by default.
```

```

positional arguments:
  DIR                project directory to create. default "."

optional arguments:
  -h, --help          show this help message and exit
  --slices SLICES     override slice count detection
  --name NAME         name of method dir and workdir, default "dev"
  --input INPUT       input directory
  --force             go ahead even though directory is not empty, or workdir
                    exists with incompatible slice count
  --tcp [HOST/PORT]  listen on TCP instead of unix sockets.
                    specify HOST (can be IP) to listen on that host
                    specify PORT to use range(PORT, PORT + 3)
                    specify both as HOST:PORT
  --no-git            don't create git repository
  --examples          copy examples to project directory

```

9.4 Accelerator Server: ax server

The Accelerator server, or daemon, needs to be running in order to execute any commands.

```
ax server
```

will start the Accelerator server, assuming that a configuration file that makes sense is in place.

```
ax server --help
```

```

usage: ax server [-h] [--keep-temp-files] [--debuggable]

optional arguments:
  -h, --help          show this help message and exit
  --keep-temp-files  keep temporary files in jobs
  --debuggable       make breakpoint() work in methods. note that this makes
                    a failing method kill the whole server.

```

Communication with the Accelerator server will take place over an UNIX socket by default. There is no need for any additional configuration to make that happen. It is possible, however, to communicate over a TCP port instead if specified in the Accelerator's configuration file.

The `debuggable`-flag will make methods execute in the same session as the server, and use that terminal as `stdin`. This way methods can be debugged using the Python Debugger. Note that a failing method will kill the server if this flag is used.

9.5 Running Build Scripts: ax run

Build scripts are executed using

```
ax run <script>
```

```
ax run --help
```

```

usage: ax run [options] [script]

positional arguments:
  script                build script to run. default "build".
                       searches under all method directories in alphabetical
                       order if it does not contain a dot.
                       prefixes build_ to last element unless specified.
                       package name suffixes are ok.

```

```
so for example "test_methods.tests" expands to
"accelerator.test_methods.build_tests".
```

optional arguments:

```
-h, --help          show this help message and exit
-f FLAGS, --flags FLAGS
                    comma separated list of flags
-q, --quick         skip method updates and checking workdirs for new jobs
-c SPEC, --concurrency SPEC
                    set max concurrency for methods, either method=N
                    or just N to set for all other methods
-w WORKDIR, --workdir WORKDIR
                    build in this workdir
                    set_workdir() and workdir= override this.
-W, --just_wait    just wait for running job, don't run any build script
-p, --full-path    print full path to jobdirs
--verbose VERBOSE  verbosity style {no, status, dots, log}
--quiet           same as --verbose=no
--horizon HORIZON time horizon - dates after this are not visible in
                    urd.latest
```

When the run command starts, it will instruct the Accelerator to scan all method directories to see if there are any new or changed methods. Thereafter, the Accelerator will proceed and scan all source workdirs to see if any new jobs have been created (by another Accelerator server). Thereafter, it will execute the build script.

9.6 Dataset Information: ax ds

The ds command gives a compact, but easy to read, overview of either

- a dataset,
- a chain of datasets, or
- available datasets in a job directory.

it provides information about column names and types, max and min values, number of rows, and balance of rows between slices.

```
usage: ax ds [options] ds [ds [...]]
```

positional arguments:

```
dataset           the job part of the dataset name can be specified in
                  the same ways as for "ax job". you can use ds~ or ds~N
                  to follow the chain N steps backwards, or ^ to follow
                  .parent. this requires specifying the ds-name, so
                  wd-1~ will not do this, but wd-1/default~ will.
```

optional arguments:

```
-h, --help          show this help message and exit
-c, --chain         list all datasets in a chain
-C, --non-empty-chain
                  list all non-empty datasets in a chain
-l, --list          list all datasets in a job with number of rows
-L, --chainedlist  list all datasets in a job with number of chained rows
-m, --suppress-minmax
                  do not print min/max column values
-n, --suppress-columns
                  do not print columns
-q, --suppress-errors
                  silently ignores bad input datasets/jobids
-s, --slices       list relative number of lines per slice in sorted
                  order
```

<code>-S, --chainedslices</code>	same as <code>-s</code> but for full chain
<code>-w, --location</code>	show where (ds/filename) each column is stored

The `dataset` option is either a *dataset*, when used with the `-s`, `-S`, and `-c` options, or a *jobid* when used with `-l` option. Datasets or jobids could be either names or absolute paths. Examples of valid datasets are `test-2`, `test-2/default`, and `/home/wdirs/test/test2/dsx`. Of these, only `test-2` is a valid jobid. Here are all options

<code>-h</code>	show help message and exit.
<code>--help</code>	
<code>-q</code>	Silently ignore any error.
<code>--quiet</code>	

When `ds` is fed with `dataset(s)`

<code>-c</code>	Print name and number of lines for all datasets in the chain.
<code>--chain</code>	
<code>-C</code>	Same as <code>-c</code> , but elide empty datasets.
<code>--non-empty-chain</code>	
<code>-m</code>	Don't print the [min, max] for each column.
<code>--suppress-minmax</code>	
<code>-n</code>	Don't print columns at all.
<code>--suppress-columns</code>	
<code>-q</code>	Bad datasets and jobids are silently ignored.
<code>--suppress-errors</code>	
<code>-s</code>	Print absolute and relative number of lines per slice for the input dataset.
<code>--slices</code>	
<code>-S</code>	Same as <code>-s</code> , but data is for the whole chain of datasets.
<code>--chain</code>	
<code>-w</code>	Print location where data for each column is stored and how it is compressed.
<code>--location</code>	

When `ds` is fed with `jobid(s)`

<code>-l</code>	Print the name and number of lines of all datasets in the input jobid.
<code>--list</code>	
<code>-L</code>	Same as <code>-l</code> , but print number of lines for the whole chain from each dataset.
<code>--chainedlist</code>	

Example invocation 1

```
ax ds test-20 -S
```

or

```
ax ds test-20/badlines
```

The argument can be one or more jobids or dataset ids. If the argument is a jobid, it is assumed that the dataset name is `default`. If there are more than one dataset in the job, a list of dataset names will be returned.

Example invocation 2

Combining `ds` will shell features can be an elegant way to extract information from a workdir. For example

```
ax ds -l -q test-{0..99}
```

will scan for datasets in the 100 first jobs of `test`. Adding the `-q` option will make `ds` suppress the warning messages for those jobs that do not contain any datasets.

Example invocation 3

Find all datasets in `test-20`

```
ax ds -l test-20
```

Example Output

```
import-2340/default
  Previous: import-2245/default
  Hashlabel: serial_number
  Columns:
    capacity_bytes      int64  [      -1, 14000519643136]
    date                date   [2019-04-01, 2019-06-30]
    failure              bool   [      False,          True]
    model                ascii
    * serial_number      ascii
  5 columns
  9,831,138 lines
  Chain length 17, from import-269 to import-2340
  Balance, lines per slice, full chain:
    1:  4.44% (6,486,330)  20:  4.35% (6,365,896)   3:  4.32% (6,323,701)
    0:  4.43% (6,472,949)  17:  4.35% (6,363,206)  19:  4.31% (6,309,295)
   11:  4.39% (6,423,397)  21:  4.35% (6,360,759)  12:  4.31% (6,306,181)
    4:  4.39% (6,421,043)  15:  4.35% (6,358,757)   8:  4.31% (6,304,311)
    6:  4.39% (6,418,617)   7:  4.34% (6,352,750)  13:  4.31% (6,303,637)
   14:  4.39% (6,417,911)  18:  4.34% (6,348,128)   5:  4.28% (6,252,735)
    2:  4.36% (6,376,139)  10:  4.34% (6,347,694)  16:  4.26% (6,230,543)
   22:  4.35% (6,366,040)   9:  4.33% (6,325,241)
  Max to average ratio: 1.020
  146,235,260 total lines in chain
```

The `max to average ratio` shows the ratio between the slice with most rows and the average number of rows. This can be interpreted as the execution time overhead for a dataset where all slices are not of the same size. When the ratio is close to one, the dataset is perfectly balanced.

9.7 Look at Specific Data in a Dataset: `ax grep`

The `grep` command is a parallel grep for datasets that is used to look at data stored in datasets or dataset chains.

```
usage: ax grep [options] [-e] pattern [...] [-d] ds [...] [[-n] column [...]]

positional arguments:
  pattern                (-e, --regexp)
  dataset                (-d, --dataset) can be specified as for "ax ds"
  columns                (-n, --column)

optional arguments:
  -h, --help            show this help message and exit
  -c, --chain           follow dataset chains
  --chain-length LENGTH, --cl LENGTH
```

```

follow chains at most this many datasets
--stop-ds DATASET      follow chains at most to this dataset
                       these options work like in ds.chain() and you can
                       specify them several times (for several datasets)
--colour [WHEN], --color [WHEN]
                       colour matched text. can be auto, never or always
-i, --ignore-case     case insensitive pattern
-v, --invert-match    select non-matching lines
-o, --only-matching   only print matching part (or columns with -l)
-l, --list-matching   only print matching datasets (or slices with -S)
                       when used with -o, only print matching columns
-H, --headers        print column names before output (and on each change)
-O, --ordered        output in order (one slice at a time)
-M, --allow-missing-columns
                       datasets are allowed to not have (some) columns
-g COLUMN, --grep COLUMN
                       grep this column only, can be specified multiple times
-s SLICE, --slice SLICE
                       grep this slice only, can be specified multiple times
-D, --show-dataset    show dataset on matching lines
-S, --show-sliceno   show sliceno on matching lines
-L, --show-lineno    show lineno (per slice) on matching lines
-f FORMAT, --format FORMAT
                       output format, csv (default) / raw / json
-t SEPARATOR, --separator SEPARATOR
                       field separator, default tab / tab-like spaces
-T LENGTH, --tab-length LENGTH
                       field alignment, always uses spaces as separator
                       specify as many as you like of TABLEN/FIELDLEN/MINLEN
                       or as NAME=VALUE (e.g. "min=3/field=24")
                       TABLEN works like normal tabs
                       FIELDLEN sets a longer minimum between fields
                       MINLEN sets a minimum len for all separators
                       use "-T/" to just activate it (sets 8/16/2)
-B NUM, --before-context NUM
                       print NUM lines of leading context
-A NUM, --after-context NUM
                       print NUM lines of trailing context
-C NUM, --context NUM
                       print NUM lines of context
                       context is only taken from the same slice of the same
                       dataset, and may intermix with output from other
                       slices. Use -O to avoid that, or -S -L to see it.

```

The pattern is a regular expression and ds are datasets. Strings and columns with special characters have to be quoted. For example

```
ax grep Alice test-0 test-1/special "first name"
```

Will look for the string Alice in the first name column of the two datasets test-0 and test-1/special.

9.8 Information about Methods: ax method

Use ax method to print a list of available methods together with an optional description.

```
ax method
```

The method description is shown by specifying a method name

```
ax method csvimport
```

Descriptions are added using the description = "some string" statement in a method.

9.9 List all build scripts: ax script

List all available build scripts using

```
ax script
```

Descriptions are added using the `description = "some string"` statement in a build script.

9.10 Information about Jobs: ax job

Get information about a job by typing

```
ax job <jobname>
```

9.11 Aliases

It is possible to create aliases for common tasks. Aliases are inserted in the file `config` in the `$XDG_CONFIG_HOME/accelerator` directory. By default, this corresponds to the file `$HOME/.config/accelerator/config`. In this file, aliases are entered in the `[alias]` section like this

```
[alias]
  path = job -P
```

This will create the alias `ax path` that is equivalent to `ax job -P`.

9.12 Show Dataset Contents: ax cat

The data in a dataset may be printed to `stdout` using the `cat` alias, for example like this

```
ax cat test-0 | less
```

This is an alias for `ax grep ""`, which greps a dataset for the empty regexp, which matches every line. A common usage is to do

```
ax cat -H0 test-0 column1 column2
```

to list a set of columns (or all, if omitted) row by row in order with a header showing label names.

If you find yourself missing the smart tab mode and colour when piping to `less` you may wish to create an alias like

```
[alias]
  catc = cat -T/ --colour=always
```

or similar.

9.13 The Urd Job Database Server: ax urd-server

By default, a local Urd server is running when the Accelerator server is running. Read more about this in section 7. A stand-alone Urd server is started by

```
ax urd-server
```

These are the options

```
ax urd-server --help
```

```
usage: ax urd-server [-h] [--path PATH] [--allow-passwordless] [--quiet]
```

options:

```
-h, --help          show this help message and exit
--path PATH         database directory (can be relative to project
                    directory) (default: urd.db)
--allow-passwordless accept any pass for users not in passwd.
--quiet            less chatty.
```

This will listen on the port/socket specified in the Accelerator's configuration file. You will probably want to configure this as `remote` so the server does not try to run another urd on the same port/socket.

9.13.1 Authorization to Urd

Authorisation to Urd could be set in the `URD_AUTH` environment variable. A common way to invoke the run command with Urd authorisation is like this

```
URD_AUTH=user:passwd ax run [script]
```

Note that the purpose of the authentication is actually *identification*. It is used to get write access to certain Urd lists. Nothing more.

9.14 Aborting a Running Job: ax abort

The command `ax abort` will abort any running job. To abort an entire build script, first type `Ctrl-C` followed by `ax abort`.

```
usage: ax abort [-h] [-q]
```

optional arguments:

```
-h, --help  show this help message and exit
-q, --quiet no output
```

9.15 Current Version: ax version

Use `ax version` to find the currently running version. The Accelerator will complain if the `ax` command version does not match the currently running server.

9.16 Workdir Information: ax workdir

By itself, `ax workdir` will print a list of all current workdirs. By specifying a workdir, all jobs in that workdir, together with execution time and information if it is *old* or *current* will be listed. All jobs in all workdirs will be listed using the `-a` option.

```
usage: ax workdir [-p] [-a | [workdir [workdir [...]]]]
```

positional arguments:

```
workdirs
```

optional arguments:

```
-h, --help          show this help message and exit
-a, --all           list all workdirs
-p, --full-path     show full path
```

9.17 The Board Web Server: ax board-server

The Board server is a web server that exposes information about jobs and results through the use of a standard web browser. This is started by default by the `ax server` command, but it can be started separately using `ax board-server`.

```
usage: ax board-server [listen_on]
runs a web server on listen_on (default localhost:8520, can be socket path)
for displaying results (result_directory)
```

Note that unlike `urd-server` this does not use the port/socket specified in the Accelerator's configuration file.

9.18 Displaying Urd Database Content: `ax urd`

```
usage: ax urd path [path [...]]
```

`path` is an optionally shortened path to an urd list.
one element is a list name, two elements are list and timestamp.
(and three elements are the whole thing, user/list/timestamp.)

use "path/since/ts" or just "path/" to list timestamps
use "/" (or nothing) to list all lists

you can also use `:urdlst:[entry]` job specifiers. `urdlst` follows the same path rules as above, `entry` is an optional argument to `joblist.get()` printing just the resultant `jobid`.

in job specifiers you can use `~` to move to forwards (later, down) along a list of timestamps and `^` to move backwards (earlier, up).

examples:

```
"ax urd example" is "ax urd eaenbrd/example/latest"
"ax urd :example:" is also "ax urd eaenbrd/example/latest"
"ax urd example/2021-04-14" is "ax urd eaenbrd/example/2021-04-14"
"ax urd :foo/bar/first:" is "ax urd foo/bar/first"
"ax urd :bar/first~:" is the second timestamp in eaenbrd/bar
"ax urd example/" is "ax urd eaenbrd/example/since/0"
```

Appendix A

Miscellaneous

DRAFT

A.1 Workdirs

Jobs are stored in *workdirs*. Workdirs are defined in the Accelerator's configuration file, where at least one workdir must be specified.

By default, the only workdir that is written to is the target workdir, while all other defined workdirs are for reading. It is possible to override this, however, by setting the `workdir=` option in the `urd.build()` call, see section 7.13.

Jobdirs are stored in the workdir by the server, and jobdirs will inherit the workdir name and add a suffix that is an incremental job counter. Here is an example of a workdir named `test`, that contains three jobdirs.

```
test/  
  .slices.conf  
  test-0/  
  test-1/  
  test-2/  
  test-LATEST -> test-2
```

The `.slices.conf` file contains the number of slices used for the workdir. The link `<workdir>-LATEST` is always pointing to the last jobdir created. This is useful for example when iteratively testing a method and accessing its data for example for plotting purposes. Each new build of the (modified) method will create a new job, and the link will always point to the most recent version.

A.1.1 Creating a Workdir

If a workdir defined in the configuration file does not exist on disk at the stated location, the server will exit and print an error stating that a directory is missing. The first time the server encounters a new directory it will initialise it in accordance with the configuration file. So, new workdirs are created by adding them to the configuration file *and* creating the corresponding directories. The Accelerator will then initiate these directories on the next startup.

The initiation process creates a file named `.slices.conf` that indicates that the directory is now a workdir. This file contains the number of slices that is used for the workdir.

A.2 Status (Progress) Reporting

During job building, it is possible to press `C-t`, i.e. `Ctrl + t` simultaneously, in the `run` shell to get status information. The built in status will report the processing state, if it is in `prepare()`, `analysis()`, or `synthesis()`. Iterators report which dataset (perhaps in a chain) that is currently being iterated, and the `blob` functions report status of file pickling.

The `status` module makes it possible to insert status reporting into any method. For more information about status reporting, see section A.3.

A.3 Generating Progress Messages: the status Module

The `status` module is used by the Accelerator to report processing state. It is also used by various functions to report iterator and file access progress. Status messages are presented in the `run` shell by pressing `C-t`, i.e. the `Ctrl` and `t` keys simultaneously.

The `status` module can be used to write progress and status messages for any function. Here is an example of how to use the `status` module

```
from accelerator import status  
...  
def analysis(sliceno):  
    msg = "reached line %d already!"  
    with status(msg % (0,)) as update:  
        for ix, data in enumerate(datasets.source.iterate(sliceno, 'data')):  
            if ix % 1000000 == 0:
```

```
update(msg % (ix,))
```

In the example above, the status message will be updated once every million iteration. By pressing `C-t` during its execution, the user will get a message telling how many lines the iterator has reached.

A.4 Working with Relative Paths

In some situations, like importing data from files, it is convenient to store the absolute path of the files as a configuration parameter and then work only with relative paths in the source code. This has two advantages.

First, it makes it possible to move input files around without forcing a re-build of the import jobs, and

second, absolute paths will not be stored in the source code.

In order to make use of relative paths, store the “system dependent” left part of the path in the Accelerator’s configuration file. There are two variables in the configuration file that can be used for this, and they have different purposes. The `input_directory` variable is intended for reading input files, and the `result_directory` is intended for writing output. See the following subsections for details.

A.4.1 The `input_directory`

Files stored in the `input_directory` can be accessed using the `CurrentJob` object either like this

```
def synthesis(job)
    fname = job.input_filename('afile.csv')
```

or like this

```
def synthesis(job)
    with job.open_input('afile.csv', 'rt') as fh:
        for line in fh:
            ...
```

Since the `input_directory` is defined in the Accelerator’s configuration file, the input directory can be moved and re-defined without having to re-execute any built jobs.

A.4.2 The `result_directory`

Storing files in job directories is great for transparency, but in some cases it is convenient to keep a reference to result files in a common place. This is the purpose of the `result_directory`. However, storing files in this directory directly would void the connection to the job that created it. A better way is to keep the file in the job directory and create a symbolic (soft) link to it in the result directory. This functionality is implemented in the `job.link_result()` function that is available to *build scripts*, and used like this

```
def main(urd):
    job = urd.build('somejob')
    job.link_result() # links the job's "result.pickle" to result_directory
    job.link_result('result.txt')
    job.link_result('result.txt', linkname='result_from_somejob.txt')
```

Note that this only works in build scripts. The `job.link_result(filename)` call will create a symbolic link named `filename` in the `result_directory` defined in the Accelerator’s configuration file. The link will point to the original file in the job directory, and will be silently replaced if it already exists.

A.5 Concurrency

By default, a job containing the `analysis()` function will be forked into `slices` parallel processes, where `slices` is specified in the Accelerator's configuration file. A standard method like `dataset_sort` will sort in parallel in all slices for maximum performance, but for large datasets and systems with little RAM, this could lead to running out of memory. A solution to this problem is to limit the number of allowed parallel processes.

This could be done either in the `build` call, using the `concurrency=` parameter, or on the command line as an option to the `ax run` command. In both cases, the limit could be set to all methods, or it could be set to a specific method only.

If `concurrency` is set to a number, like this

```
concurrency=3
```

the number of parallel processes is limited to this number for all methods. Alternatively, `concurrency` can be specified for a single method like this

```
concurrency="dataset_sort=3"
```

Appendix B

Setup and Installation

This chapter covers how to install the Accelerator, how to configure it, and how to set up a new project.

DRAFT

B.1 Install the Accelerator

B.1.1 Using the pip command

The easiest way to install the Accelerator is by fetching it from the PyPi repository

```
pip install accelerator
```

Some prefer to install to a virtual environment and do something in line with the following

```
python3 -m venv accvenv
source accvenv/bin/activate
pip install accelerator
```

This will install the Accelerator to the `accvenv` virtual environment. Now, use for example the following command

```
ax --help
```

to check that the installation worked. The next step is to set up a project.

B.2 Set up a New Project

In order to run, the Accelerator needs to have these things in place

- at least one *workdir* to store data in,
- most likely a new *method package* directory to store new code in, and
- a *configuration file* to set things up.

This is all taken care of by the `init` command.

```
ax init
```

For more information see section 9.3. A typical project setup will look like this

```
myproject/
  accelerator.conf
  dev/
    (methods.conf)
    a_method.py
    build.py
```

where methods are stored in the `dev` directory. `methods.conf` is only required when not using `auto-discover`, see `sectionsec:configfile`.

B.3 Run the Tests

A rather extensive test suite is included in the Accelerator installation. To run this, enable the test package in the configuration file:

```
method packages:
  ...
  accelerator.test_methods
```

start the server using

```
ax server
```

and in another shell start the test

```
ax run tests
```

Since the tests include testing of different character encodings, you may end up with a

```
Exception: Failed to enable numeric_comma, please install at least one
of the following locales: da_DK nb_NO nn_NO sv_SE fi_FI en_ZA es_ES
es_MX fr_FR ru_RU de_DE nl_NL it_IT
```

On a Debian-based machine, locales can be configured using

```
dpkg-reconfigure locales
```

which has to be run with root privileges.

B.4 Server Configuration File

The configuration file specifies which method packages and workdirs that are available for a project. A template configuration file can be generated using the `init` command as described in section 9.3. Below is an example of a configuration file.

```
# The configuration is a collection of key value pairs.
#
# Values are specified as
# key: value
# or for several values
# key:
#     value 1
#     value 2
#     ...
# (any leading whitespace is ok)
#
# Use ${VAR} or ${VAR=DEFAULT} to use environment variables.

slices: 23
workdirs:
    test /zbd/workdirs/test
    import ${HOME}/workdirs/import
    live wdirs/live

# Target workdir defaults to the first workdir, but you can override it.
# (this is where jobs without a workdir override are built)
#target workdir: dev

method packages:
    dev auto-discover
    # accelerator.examples auto-discover
    accelerator.standard_methods
    # accelerator.test_methods

# listen directives can be [host]:port or socket path.
# urd should be prefixed with "local" to run it together with the server
# or "remote" to not run it together with the server.
listen: .socket.dir/server
board listen: .socket.dir/board
urd: local .socket.dir/urd

result directory: ${HOME}/accelerator/results
input directory: /zbd/data/backblaze

# If you want to run methods on different python interpreters you can
# specify names for other interpreters here, and put that name after
# the method in methods.conf.
# You automatically get four names for the interpreter that started
# the server: DEFAULT, 3, 3.7 and 3.7.3 (adjusted to the actual
```

```
# version used). You can override these here, except DEFAULT.
# interpreters:
#     2.7 /path/to/python2.7
#     test /path/to/beta/python
```

The configuration file above specifies 23 slices and three workdirs, called `test`, `import`, and `live`. The `test` workdir is specified using an absolute path, the `import` workdir is specified relative to the user's home directory using the shell environment variable `$HOME`, and the `live` workdir is specified using a path relative to the location of the configuration file itself.

The workdir that is specified first is the *target workdir*, where jobs are written to by default. All other specified workdirs will by default only be used for reading. Any of the workdirs specified could be written to, though, using the `set_workdir=` option to the `build` command, as described on page 7.14.

Methods packages available for use are the `standard_methods` bundled with the Accelerator, and methods defined in the directory `dev` (if defined in `dev/methods.conf`).

name	description
<code>slices</code>	Number of slices used for the project.
<code>workdirs</code>	<p>A list of paths to workdir directories. At least one workdir needs to be defined. All workdirs that are used together must have the same number of slices. It is possible to use shell environment variables such as <code>\${HOME}</code> when specifying workdirs. Path starting with a slash (<code>/</code>) are absolute paths, all other paths are relative to the location of the configuration file itself.</p> <p>Unless overridden by the <i>target workdir</i>, the first workdir in the list will be the default <i>target workdir</i> that is used for all writing. Other specified workdirs will only be read from, unless overridden by the <code>build</code> call as described on page 7.14.</p>
<code>target workdir</code>	Name of the <i>target workdir</i> . If specified this overrides the first item in the <code>workdirs</code> list.
<code>method packages</code>	<p>A list of directories containing methods. These will be the only directories where the Accelerator can “see” methods. <code>standard_methods</code> is bundled with the Accelerator and is commonly used.</p> <p>A package specified alone on a row indicates that its corresponding directory contains a file named <code>methods.conf</code> that specify which methods in the package that can be executed.</p> <p>If the package is specified with the string <code>auto-discover</code>, like in the example above, any method in the package can be executed by the Accelerator, and no <code>methods.conf</code> file is required. It can still be present in order to specify which interpreters to use on a method level though.</p>
<code>urd</code>	If present, an URL to the Urd server.
<code>result directory</code>	A common path that is available to all jobs. Use the <code>job.link_result()</code> -function to create symbolic links from files in job directories to this directory. Just like workdirs, this path is either absolute or relative to the location of the configuration file.

<code>input directory</code>	Default root path for <code>csvimport</code> . This is to avoid rebuilds of imports if input files are moved to another directory. (This typically happens when setting up a similar system on another physical machine.) See section A.4.1 on how to get access to <code>input_directory</code> from any method. Just like <code>workdirs</code> , this path is either absolute or relative to the location of the configuration file.
<code>interpreters</code>	Name and path to python executables. These are used in <code>methods.conf</code> to specify specific Python versions (or virtual environments) for individual methods. If unspecified, methods will be executed using the same binary that runs the Accelerator's server process.

It is possible to assign values in the configuration file using shell environment variables. In the example above, `workdirs` are specified relative to `${HOME}`, for example. In general, the assignment is `${VAR=DEFAULT}`.

B.5 Setting up a Standalone Urd-server

The main server program will start a local Urd server by default. This server is for local use only. If the Urd server should be used to share information between users, a standalone server needs to be set up. This section shows how to do that.

The Urd database is stored in a directory of its own. If the project has been set up using `ax init`, there is already an `urd.db/` directory created. Otherwise it has to be created. In either case, the configuration file should be modified so that it reads

```
urd: remote localhost:<port>
```

where `<port>` could be any available port, such as for example 30025. The next step is to set up passwords as described in section B.5.2.

B.5.1 Starting Urd

The Urd server daemon is started using

```
ax urd-server
```

and as always, add `-help` for more information. A common flag is `-allow-passwordless` that allows entries to be written without using any passwords.

B.5.2 The passwd file

The `passwd` file in the `urd.db/` directory stores write access authentication. The file format is straightforward, each line is a user–password pair as follows

```
user:password
```

For example, if the file contains the following line

```
ab:secret
```

A build script run like this

```
URD_AUTH=ab:secret ax run script
```

will have write access to all lists belonging to the user `ab`, such as for example the `ab/test` and `ab/import` lists. But it can not write to lists belonging to other users, such as `cd/import`. It can always read all lists, though.

B.5.3 Connecting to a Remote Urd Server

In order to connect to a remote Urd server, the configuration file should be modified like this

```
urd: remote localhost:<port>
```

i.e. exactly the same as for the Urd server itself. An accelerator server with this configuration file will connect to and use the external Urd server for all job list storage and retrieval.

B.5.4 The Urd Database

The Urd database has the following structure

```
database_root/  
  passwd  
  database/  
    user1/  
      list1  
      list2  
    user2/  
      list3
```

The database is a transaction log database, meaning that data is *never* modified or deleted. Information is only added. Also, all data in the database is in a human readable format.

B.6 Setting up a Standalone Board-server

The main server program will start a local board server by default. This behaviour can be disabled in the Accelerator's configuration file. In that case, to start the board server manually, use the

```
ax board
```

command.

Appendix C

Classes

The Accelerator is programmed using an object oriented approach. This chapter outlines the most common classes and its member functions.

DRAFT

C.1 The Job and CurrentJob Classes

The `Job` and `CurrentJob` classes are similar, but used in different contexts:

The `Job` class is used to represent and operate on *existing* jobs. An object of this class is returned from `job build()` calls as well as when retrieving jobs from `Urd` or a `JobList` object.

The `CurrentJob` class is an extension that provides mechanisms for operations performed while a job is *executing*, such as saving files to the job's `jobdir`.

The classes are derived from the `str` class, and objects of these classes decay to (unicode) strings when pickled. The following attributes are available on both the `Job` and `CurrentJob` classes:

name	description
<code>chain()</code>	Job chain
<code>dataset()</code>	Return a named dataset from the job.
<code>datasets</code>	List of datasets in job.
<code>filename()</code>	Return absolute path to a file in a job.
<code>files()</code>	Return list of files created by job.
<code>json_load()</code>	Load a json file from the job's directory.
<code>link_result()</code>	Create a soft link from a file in a job's directory to <code>result_directory</code> .
<code>load()</code>	Load a pickle file from the job's directory.
<code>method</code>	The job's method. This can be overridden by <code>name=</code> if job instance is output from <code>Urd</code> or a <code>build()</code> call.
<code>number</code>	The job number as an <code>int</code> .
<code>open()</code>	Similar to standard <code>open</code> , use to open files.
<code>output()</code>	Return what the job printed to <code>stdout</code> and <code>stderr</code> .
<code>params</code>	Return a dict corresponding to the file <code>setup.json</code> for this job.
<code>path</code>	The filesystem directory where the job is stored.
<code>post</code>	Return a dict corresponding to the job's <code>post.json</code> .
<code>withfile()</code>	A <code>JobWithFile</code> with this job.
<code>workdir</code>	The <code>workdir</code> name (the part before <code>-number</code> in the <code>jobid</code>).

In addition, the `CurrentJob` class has these unique attributes:

name	description
<code>datasetwriter()</code>	Returns a <code>DatasetWriter</code> object. See documentation for <code>DatasetWriter()</code> , section C.6.
<code>input_filename()</code>	Full filename of file in <code>input_directory</code> .
<code>json_save()</code>	Store a json file in the current job directory.
<code>open_input()</code>	Open a file in <code>input_directory</code> .
<code>register_file()</code>	Record a file produced by this job.
<code>save()</code>	Store a pickle file in the current job directory.

Detailed description of the functions, where necessary, follows.

C.1.1 Job.chain()

This works like `dataset.chain()`, but for chains of jobs.

name	default	description
length	-1	Length of chain.
reverse	False	Reverse direction of chain.
stop_job	None	Chain to here.

C.1.2 Job.dataset()

name	default	description
name	default	

Get a dataset instance from a job.

C.1.3 Job.filename()

name	default	description
filename	Mandatory	Name of file in job directory.
sliceno	None	Set to current slice number if sliced, otherwise None.

Return the absolute (full path) filename to a file stored in the job. If the file is sliced, a particular slice file can be retrieved using the `sliceno` parameter. Sliced files are described in section 4.7.2.

C.1.4 Job.files()

name	default	description
pattern	*	Return only files matching pattern

This method returns a list of all filenames corresponding to files created by the job using the functions `CurrentJob.open()`, `CurrentJob.save()`, `CurrentJob.json_save()` or registered using `CurrentJob.register_file()`. The list can be filtered using the `pattern` option. Filtering is based on Python's `fnmatch`-functionality.

C.1.5 Job.json_load()

name	default	description
filename	result.json	Name of file.
sliceno	None	

Load a file from a job, in JSON format.

C.1.6 Job.link_result()

name	default	description
filename	<i>Mandatory</i>	Name of file in job directory.
linkname	<i>None</i>	Name of link if set.

Use to create a soft link from a file in a job directory to the `result_directory`.

NOTE: This only works for `Job` instances, and not `CurrentJob` instances. This is for reproducibility reasons. Links in `result_directory` cannot be recreated if created in a job, since jobs can only be executed once.

C.1.7 `Job.load()`

name	default	description
filename	<code>result.pickle</code>	Name of file.
sliceno	<i>None</i>	
encoding	<code>bytes</code>	

Load a file from a job in Python's pickle format.

C.1.8 `Job.open()`

name	default	description
filename	<i>Mandatory</i>	Name of file.
mode	<code>r</code>	Open file in this mode, see Python's <code>open()</code>
sliceno	<i>None</i>	Read or write sliced files.
encoding	<i>None</i>	Same as Python's <code>open()</code>
errors	<i>None</i>	Same as Python's <code>open()</code>
temp	<i>None</i>	Control file persistence. See text.

This is a wrapper around the standard `open` function with some extra features. Note that

- `Job.open()` can only read files, not write them, and therefore “`r`” flag must be set.
- `CurrentJob.open()` can both read and write.
- `CurrentJob.open()` must be used as a context manager, like this

```
with job.open(...) as fh:
    ....
```

- `CurrentJob.open()` can use the `temp` flag to modify the persistence of written files.
- `CurrentJob.open()` will register the file to the current job.

The `temp` argument is used to control the persistence of files written using `.open()`. This is useful mainly for debug purposes, and explained in section C.1.17. Sliced files are described in section C.1.16.

C.1.9 `Job.output()`

name	default	description
what	<i>None</i>	Which functions to return output from.

Get everything a job has printed to `stdout` and `stderr` in a string variable. The parameter `what` can be set to

- `None`, which returns everything,
- `prepare`, which returns everything from `prepare`,
- `analysis`, which returns everything from `analysis`,
- `synthesis`, which returns everything from `synthesis`, or
- a number, which returns output from the corresponding `analysis` slice.

C.1.10 `Job.withfile()`

name	default	description
<code>filename</code>	<i>Mandatory</i>	Name of file.
<code>sliced</code>	<code>False</code>	Boolean indicating if the file is sliced or not.
<code>extra</code>	<code>None</code>	Any additional information to the job to be built.

The `.withfile()` is used to highlight a specific file in a job and feed it to another job `build()`. The file could be sliced.

C.1.11 `CurrentJob.input_filename()`

Return the full filename of a file stored in the `input_directory`.

name	default	description
<code>filename</code>	<i>Mandatory</i>	Name of file.

C.1.12 `CurrentJob.json_save()`

name	default	description
<code>obj</code>	<i>Mandatory</i>	
<code>filename</code>	<code>result.json</code>	
<code>sliceno</code>	<code>None</code>	
<code>sort_keys</code>	<code>True</code>	
<code>temp</code>	<code>None</code>	

For `CurrentJob` instances only. Save data into the current job's directory in JSON format. The `temp` argument is used to control the persistence of files written using `.json_save()`. This is useful mainly for debug purposes, and explained in section C.1.17.

C.1.13 `CurrentJob.open_input()`

Open a file in the `input_directory`.

name	default	description
<code>filename</code>	<i>Mandatory</i>	Name of file.
<code>mode</code>	<code>r</code>	Open file in this mode, see Python's <code>open()</code>
<code>encoding</code>	<code>None</code>	Same as Python's <code>open()</code>

errors	None	Same as Python's <code>open()</code>
--------	------	--------------------------------------

C.1.14 `CurrentJob.register_file()`

Register a file to the running job. This happens automatically when using `Job.open()`, but if the file was produced in a way where that is not practical this can be used to manually register the file.

name	default	description
filename	Mandatory	Name of file.

C.1.15 `CurrentJob.save()`

name	default	description
obj	Mandatory	
filename	result.pickle	
sliceno	None	
temp	None	

For `CurrentJob` instances only. Save data into the current job's directory in Python's pickle format. The `temp` argument is used to control the persistence of files written using `.save()`. This is useful mainly for debug purposes, and explained in section C.1.17.

C.1.16 Sliced Files

A *sliced* file is actually a set of files used to store data independently in each `analysis()` process using a common name. The functions that operate on files, such as for example `.open()` and `.load()`, can switch to sliced files using the `sliceno` parameter. From a user's perspective, they always appear to work on single files. For example

```
def analysis(sliceno, job):
    data = ...
    job.save(data, "mydata", sliceno=sliceno, temp=False)
```

will create a set of files `mydata.%d`, where `%d` is replaced by the slice number. In this way, data can be passed "in parallel" between different jobs.

C.1.17 File Persistence

The `temp` argument controls persistence of files stored using `.open()`, `.save()`, or `.json_save()`. By default it is being set to `False`, which implies that the stored file is *not* temporary. But setting it to `True`, like in the following

```
job.save(data, filename, temp=True)
```

will cause the stored file to be deleted upon job completion. The functionality can be combined with the *debug* mode, see below.

temp	"normal" mode	debug mode
False	stored	stored
True	stored and removed	stored

Debug mode is active if the Accelerator server is started with the `--debug` flag.

C.2 The JobWithFile Class

The `JobWithFile` class is used to create a job input parameter from a file stored in a job.

name	description
<code>filename()</code>	Return filename.
<code>load()</code>	Load file contents. Takes an <code>encoding=-</code> parameter.
<code>json_load()</code>	Load JSON file contents.
<code>open()</code>	Returns a filehandle to the open file. Takes <code>encoding=-</code> and <code>errors=-</code> parameters.

All above functions take the argument `sliceno`, which default is set to `None`, indicating that it is actually a single file on disk. If `sliceno` is set, it is assumed that the file is sliced, see section C.1.16, and the function will look up that slice of the file only.

C.3 The JobList Class

Objects of the `JobList` class are returned by member functions to the `Urd` class. They are used to group sessions of jobs together.

name	description
<code>find()</code>	Return a new <code>JobList</code> with only jobs with that method or name in it.
<code>get()</code>	Return the latest <code>Job</code> with that method or name.
<code>[<method>]</code>	Same as <code>.get</code> but error if no job with that method or name is in the list.
<code>as_tuples</code>	The <code>JobList</code> represented as <code>(method, jid)</code> tuples.
<code>pretty</code>	Return a prettified string version of the <code>JobList</code> .
<code>exectime</code>	Execution times in total as well as per method.
<code>print_exectimes()</code>	Print execution time information to <code>stdout</code> .

Detailed description of the functions, where necessary, follows.

C.3.1 `JobList.find()`

name	default	description
<code>method</code>	<i>Mandatory</i>	Method or name to find.

Return a new `JobList` with all jobs in the current `JobList` matching the `method` argument. The matching part is either the unique name of the method's source code, or the name optionally given at build time using the `name=` argument.

C.3.2 `JobList.get()`

name	default	description
<code>method</code>	<i>Mandatory</i>	Method or name to find.
<code>default</code>	<i>None</i>	Return the latest matching job.

Return the latest job that matches the `method` argument. The matching part is either the unique name of the method's source code, or the name optionally given at build time using the `name=` argument. If no matches are found, it will return the `default` argument.

C.3.3 `JobList.print_exectimes()`

name	default	description
<code>verbose</code>	<i>True</i>	In addition to total time, print execution time for each method in list.

Print total execution time for the `JobList`, and, conditionally, execution time for each job in the list, to `stdout`.

C.4 The Dataset Class

The `Dataset` class is used to operate on small or large datasets stored on disk. It decays to a (unicode) string when pickled.

name	description
<code>min()</code>	Min value for a specified column over the whole chain.
<code>max()</code>	Max value for a specified column over the whole chain.
<code>columns</code>	A <code>dict</code> from column to properties, such as type, min, and max values.
<code>previous</code>	The dataset's previous dataset, if it exists, <code>None</code> otherwise.
<code>parent</code>	The dataset's parent dataset, if it exists, <code>None</code> otherwise.
<code>filename</code>	The dataset's filename, if it exists. (<code>csvimport</code> sets this.)
<code>hashlabel</code>	Column used for hash partitioning, or <code>None</code> .
<code>caption</code>	The dataset's caption.
<code>lines</code>	A <code>list</code> with number of lines per slice.
<code>shape</code>	A tuple containing number of columns and number of lines in dataset.
<code>link_to_here()</code>	Used to associate a subjob's dataset with the current job, see section 4.10.
<code>merge()</code>	Merge this dataset with another dataset, see section C.4.3.
<code>chain()</code>	A <code>DatasetChain</code> object, see section C.5
<code>iterate_chain()</code>	Iterator over chains, see chapter 6.
<code>iterate()</code>	Iterator over dataset see chapter 6.
<code>iterate_list()</code>	Iterator over a list of datasets, see chapter 6.

Detailed description of the functions, where necessary, follows.

C.4.1 `Dataset.min()`, `Dataset.max()`

name	default	description
<code>column</code>	<i>Mandatory</i>	Min/max value of column

Minimum or maximum value for column over the dataset.

C.4.2 `Dataset.link_to_here()`

name	default	description
<code>name</code>	<code>default</code>	The new name of the dataset.
<code>column_filter</code>	<code>None</code>	Iterable of columns to include, or <code>None</code> to get all.
<code>override_previous</code>	<code>_no_override</code>	Set this to the new previous.

Use this to expose a subjob as a dataset in your job, like in this example:

```
def synthesis():
    job = build('ex')
    job.dataset().link_to_here(name='new')
```

The current job will now appear to have a dataset named `new`, that is actually a link to the subjob's `default` dataset. It is possible to filter which columns should be visible in the link using `column_filter`. For chaining purposes, it is possible for the link to expose a parent dataset of choice, set using the `override_previous` parameter.

C.4.3 Dataset.merge()

name	default	description
<code>other</code>	<i>Mandatory</i>	Merge with this dataset.
<code>name</code>	"default"	Name of new dataset
<code>previous</code>	<code>None</code>	The new dataset's <code>previous</code> dataset.
<code>allow_unrelated</code>	<code>False</code>	Set this if the datasets do not share a common ancestor.

Merge this and other dataset. Columns from the other dataset take priority. If datasets do not have a common ancestor you get an error unless `allow_unrelated` is set. The new dataset always has the `previous` specified here (even if `None`). Returns the new dataset.

C.4.4 Dataset.chain()

name	default	description
<code>length</code>	-1	Number of datasets in chain. The default value of -1 will include all datasets in chain.
<code>reverse</code>	<code>False</code>	Reverse order of chain.
<code>stop_ds</code>	<code>None</code>	If set, chain will start at the dataset after <code>stop_ds</code> .

This function will return a `DatasetChain` object, see section C.5.

C.5 The DatasetChain Class

These are lists of datasets returned from `Dataset.chain`. They exist to provide some convenience methods on chains.

name	description
<code>min()</code>	Min value for a specified column over the whole chain.
<code>max()</code>	Max value for a specified column over the whole chain.
<code>lines()</code>	Number of rows in this chain, optionally for a specific slice.
<code>column_counts()</code>	The number of datasets each column appears in.
<code>column_count()</code>	Number of datasets in this chain that contain a specified column.
<code>with_column()</code>	Return a new chain without any datasets that don't contain a specified column.
<code>none_support()</code>	Return <code>True</code> if any dataset in the chain has <code>None</code> -support for this column.
<code>iterate(...)</code>	Same arguments as <code>Dataset.iterate()</code> . Will iterate over the whole chain.

Detailed description of the functions, where necessary, follows.

C.5.1 `DatasetChain.min()`, `DatasetChain.max()`

name	default	description
<code>column</code>	<i>Mandatory</i>	Min/max value of column, see text.

Minimum or maximum value for column over the whole chain. Will be `None` if no dataset in the chain contains `column`, if all datasets are empty or if `column` has a type without min/max tracking.

C.5.2 `DatasetChain.lines()`

name	default	description
<code>sliceno</code>	<code>None</code>	If set, return number of lines in specified slice.

Number of rows in this chain, optionally for a specific slice.

C.5.3 `DatasetChain.column_counts()`

Return a Python `Counter`, `{colname: occurrences}`, holding the number of datasets each column appears in. Takes no options.

C.5.4 `DatasetChain.column_count()`

name	default	description
<code>column</code>	<i>Mandatory</i>	A column name.

Number of datasets in this chain that contain a specified column.

C.5.5 DatasetChain.with_column()

name	default	description
<code>column</code>	<i>Mandatory</i>	A column name.

Return a new `DatasetChain` with all datasets in this chain containing a specified column.

DRAFT

C.6 The DatasetWriter Class

The `DatasetWriter` class is used to create datasets. Datasets could be stand-alone, part of a chain, or an extension (new columns) to an existing dataset.

The class has a number of member functions, described below, that may be used for dataset creation. Alternatively, the new dataset could be set up using the `DatasetWriter` constructor. The constructor approach is currently only documented in the source code, see `dataset.py`.

name	description
<code>add()</code>	Add a new column to the dataset under creation.
<code>hashcheck()</code>	Check if value belongs in current slice.
<code>set_slice()</code>	Set which slice that will receive the next write.
<code>enable_hash_discard()</code>	Make the write functions silently discard data that does not hash to the current slice.
<code>get_split_write()</code>	Get a writer object, see section 5.9.2.
<code>get_split_write_list()</code>	Get a writer object, see section 5.9.2.
<code>get_split_write_dict()</code>	Get a writer object, see section 5.9.2.
<code>discard()</code>	Discard the dataset under creation.
<code>finish()</code>	Call this if dataset is to be used before creating job finishes, e.g. if the dataset under creation is input to a subjob.

Detailed description of the functions, where necessary, follows.

C.6.1 `DatasetWriter.add()`

name	default	description
<code>colname</code>	<i>Mandatory</i>	Name of new column.
<code>coltype</code>	<i>Mandatory</i>	Type of new column.
<code>none_support</code>	False	Set to True to allow storing Nones .

Add a new column to a dataset in creation. This example will create an `age` column of type `number`, where the values could also be `None`.

```
dw.add('age', 'number', none_support=True)
```

All dataset types are described in chapter 5.

C.6.2 `DatasetWriter.hashcheck()`

name	default	description
<code>value</code>	<i>Mandatory</i>	Some data.

Check if a value belongs to the current slice. Return **True** if `value` belongs to the current slice, **False** otherwise.

C.6.3 `DatasetWriter.set_slice()`

name	default	description
<code>sliceno</code>	<i>Mandatory</i>	Slice number to use for writing.

Specify which slice that will receive the next `write(s)`. Use this if writing data in `prepare()` or `synthesis()`.

C.6.4 `DatasetWriter.enable_hash_discard()`

Takes no options. Set this in each slice or after each `set_slice()` to make the writer discard values that do not belong to the current slice.

DRAFT

C.7 The Urd Class

name	description
<code>get()</code>	Get an Urd item from a specified list and timestamp.
<code>latest()</code>	Get the latest Urd item for a specified list.
<code>first()</code>	Get the first Urd item for a specified list.
<code>peek()</code>	Get an Urd item from a specified list and timestamp without recording.
<code>peek_latest()</code>	Get the latest Urd item for a specified list without recording.
<code>peek_first()</code>	Get the first Urd item for a specified list without recording.
<code>since()</code>	Get all timestamps later than a specified timestamp for a specified list.
<code>list()</code>	List all Urd lists
<code>begin()</code>	Start a new Urd session.
<code>abort()</code>	Abort a running Urd session.
<code>finish()</code>	Finish a running Urd session and store its contents.
<code>truncate()</code>	Discard all Urd items later than a specified timestamp for a specified list.
<code>set_workdir()</code>	Set the target workdir.
<code>build()</code>	Build a job.
<code>build_chained()</code>	Build a job with chaining.
<code>warn()</code>	Add a warning message to be displayed at the end of the build.

Detailed description of the functions, where necessary, follows.

C.7.1 `Urd.get()`

name	default	description
<code>path</code>	<i>Mandatory</i>	
<code>timestamp</code>	<i>Mandatory</i>	

Get an Urd item with specified list and timestamp. The operation is recorded in the current Urd session.

C.7.2 `Urd.latest()`

name	default	description
<code>path</code>	<i>Mandatory</i>	

Get the latest job in a specified Urd list. The operation is recorded in the current Urd session.

C.7.3 `Urd.first()`

name	default	description
------	---------	-------------

<code>path</code>	<i>Mandatory</i>
-------------------	------------------

Get the first job in a specified Urd list. The operation is recorded in the current Urd session.

C.7.4 `Urd.peek()`

name	default	description
<code>path</code>	<i>Mandatory</i>	
<code>timestamp</code>	<i>Mandatory</i>	

Same as `.get()`, but without recording the dependency.

C.7.5 `Urd.peek_latest()`

name	default	description
<code>path</code>	<i>Mandatory</i>	

Same as `.latest()`, but without recording the dependency.

C.7.6 `Urd.peek_first()`

name	default	description
<code>path</code>	<i>Mandatory</i>	

Same as `.first()`, but without recording the dependency.

C.7.7 `Urd.since()`

name	default	description
<code>path</code>	<i>Mandatory</i>	
<code>timestamp</code>	<i>Mandatory</i>	

Return a list of all timestamps more recent than the input `timestamp` for a specified Urd list.

C.7.8 `Urd.list()`

Return a list of all available Urd lists.

C.7.9 `Urd.begin()`

name	default	description
<code>path</code>	<i>Mandatory</i>	
<code>timestamp</code>	<i>Mandatory</i>	
<code>caption</code>	<i>None</i>	
<code>update</code>	<i>False</i>	

Start a new Urd session.

C.7.10 `Urd.abort()`

Abort the current Urd session, discard its contents.

C.7.11 `Urd.finish()`

name	default	description
path	<i>Mandatory</i>	
timestamp	<i>Mandatory</i>	
caption	<i>None</i>	

Finish the current Urd session and store it in the Urd database.

C.7.12 `Urd.truncate()`

name	default	description
path	<i>Mandatory</i>	
timestamp	<i>Mandatory</i>	

Discard everything later than `timestamp` for the specified Urd list.

C.7.13 `Urd.set_workdir()`

name	default	description
workdir	<i>Mandatory</i>	

Set target workdir. It can be set to any workdir present in the Accelerator's configuration file.

C.7.14 `Urd.build()`

name	default	description
method	<i>Mandatory</i>	Method to build.
options	<code>{}</code>	Input options.
datasets	<code>{}</code>	Input datasets.
jobs	<code>{}</code>	Input jobs.
name	<i>None</i>	Record job using this name instead of method name.
caption	<i>None</i>	Optional caption
workdir	<i>None</i>	Store job in this workdir.

Build a job. If an Urd session is running, the job and its dependencies will be recorded.

C.7.15 `Urd.build_chained()`

Build a chained job. Same options as `.build()`. See chapter 5 for more information.

C.7.16 Urd.warn()

Print a string to `stdout` when the build script ends with no errors.

name	default	description
<code>line</code>	<i>Mandatory</i>	Some string

DRAFT