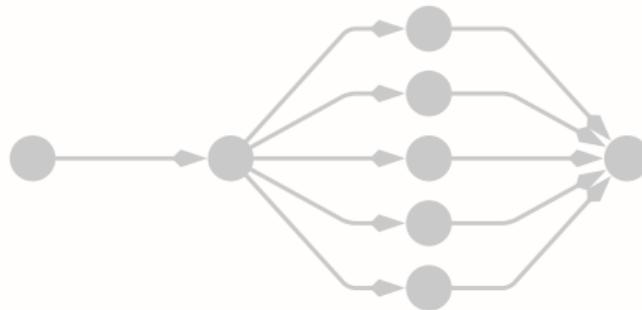


# GNU PARALLEL

---

Parallelizing and  
Distributing programs with the Shell

Felix Rieg and Florian Sihler



July 3, 2022  
CCPDP, Ulm University

**2.1**

# Motivation

[1]: *GNU Parallel 20210822*  
("Kabul")  
Tange, 2021



**Motivation**



Background



GNU parallel



Inner Workings



Outlook



## 2.2

# Motivation

Doing stuff parallel.

## 2.3

# Motivation

Doing stuff parallel.

With the commandline!

**3.1**

## Pipes

# Unix Pipelines



**3.2**

## Pipes

# Unix Pipelines



**3.3**

## Pipes

# Unix Pipelines



**3.4**

## Pipes

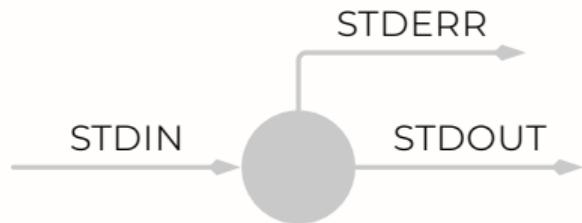
# Unix Pipelines



3.5

## Pipes

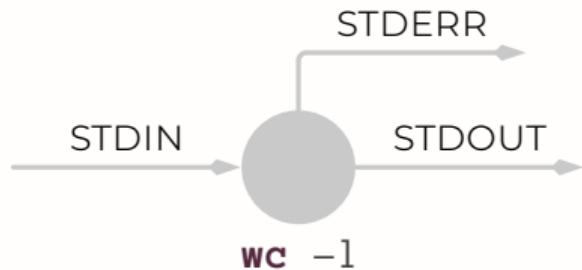
# Unix Pipelines



3.6

## Pipes

# Unix Pipelines



3.7

## Pipes

# Unix Pipelines



### 3.8

## Pipes

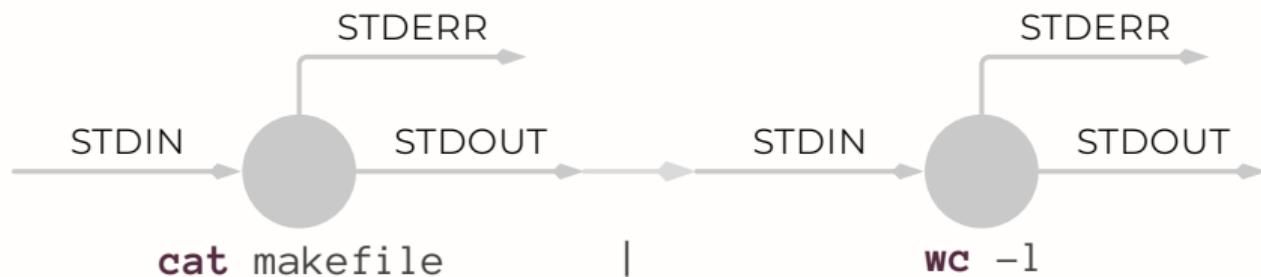
## Unix Pipelines



3.9

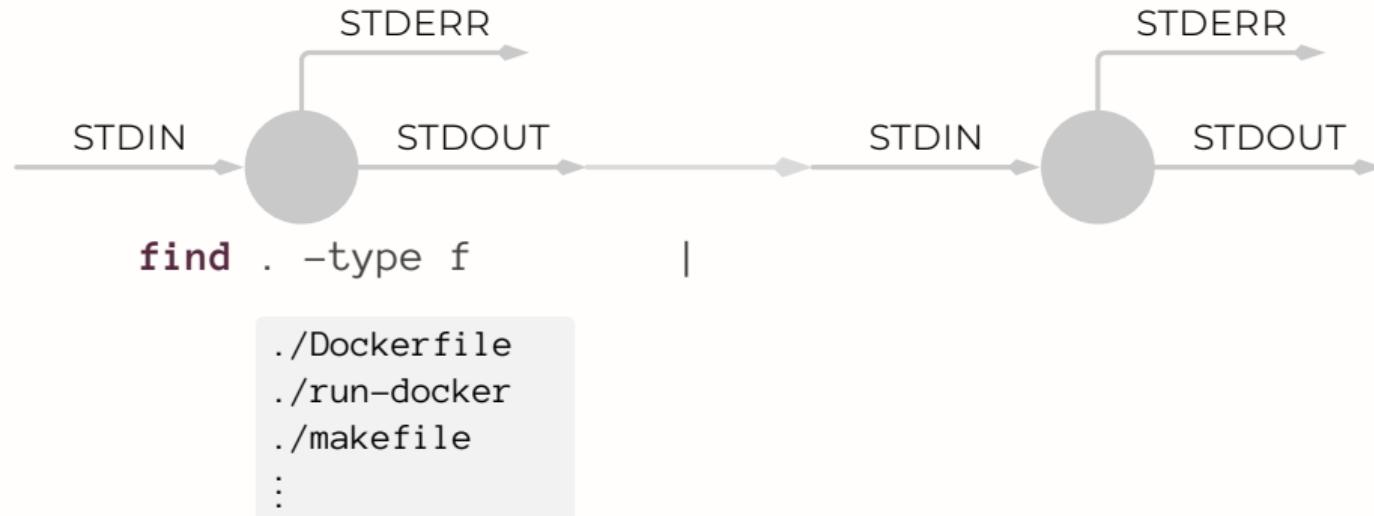
## Pipes

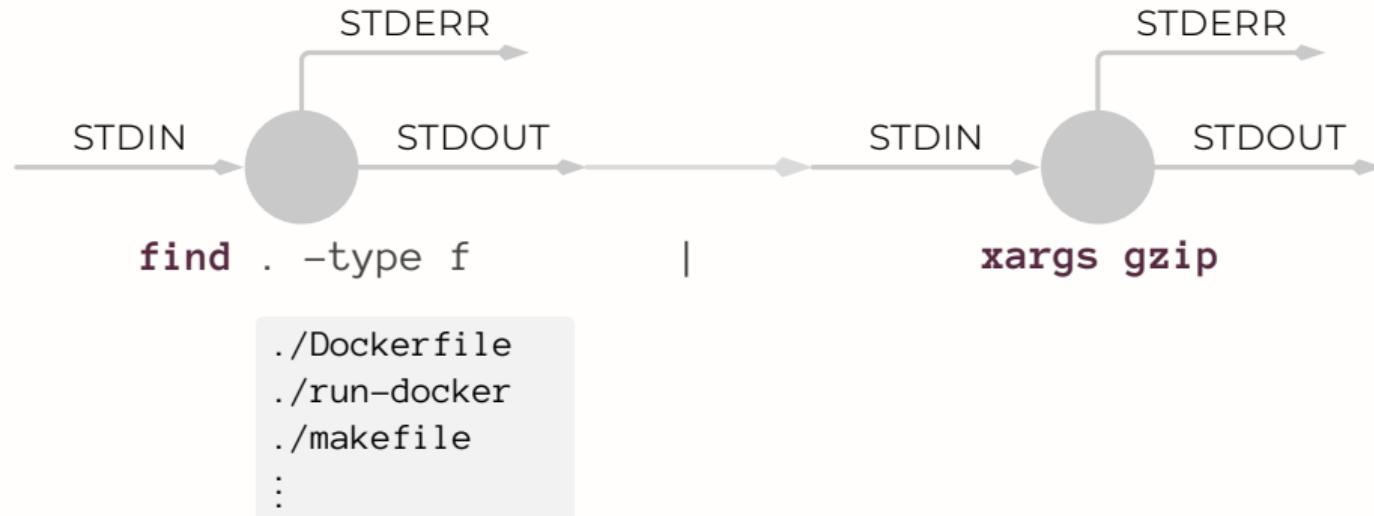
# Unix Pipelines

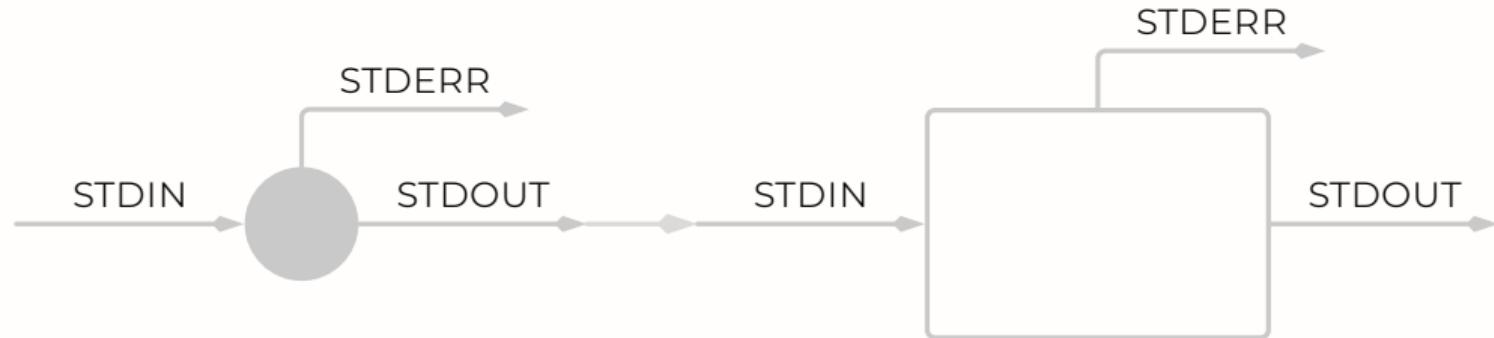










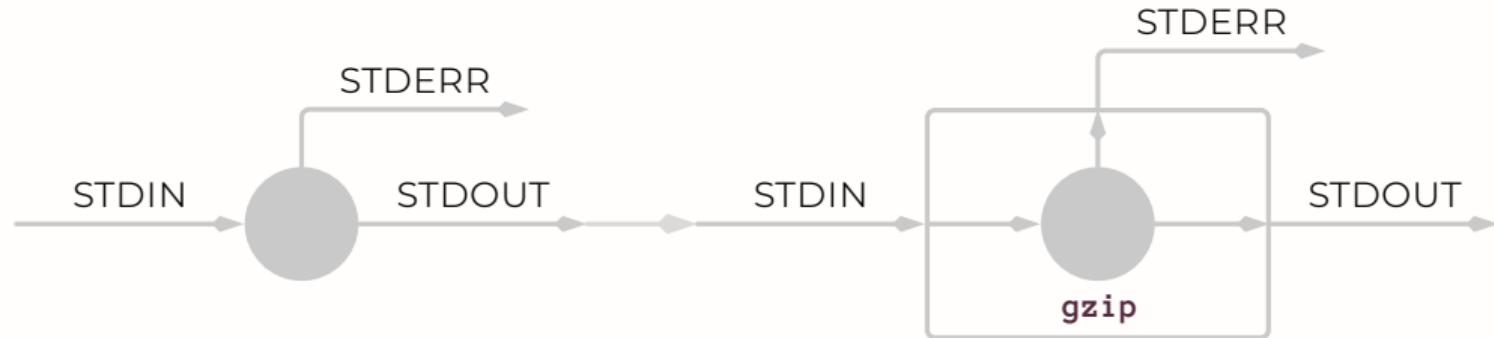


`find . -type f`

|

`xargs gzip`

```
./Dockerfile  
./run-docker  
./makefile  
⋮
```

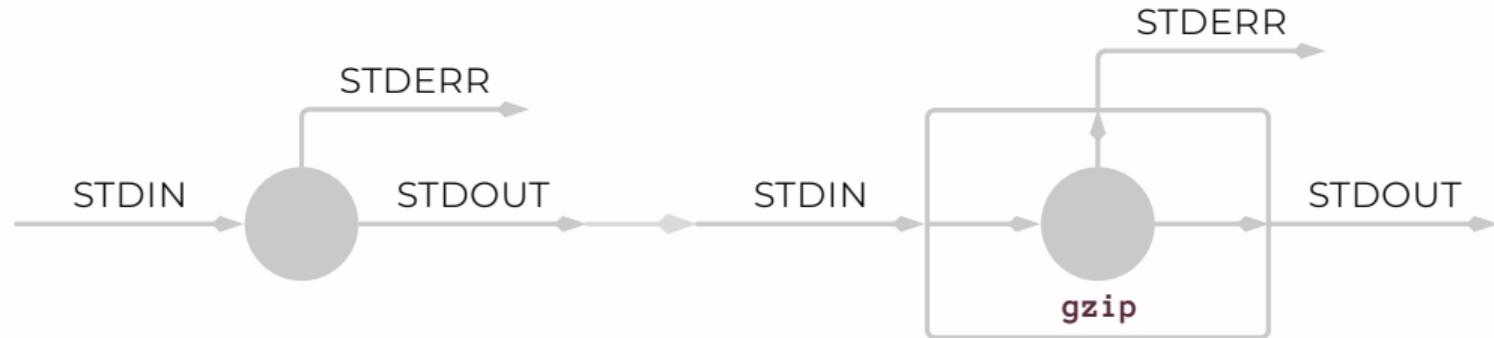


`find . -type f`

|

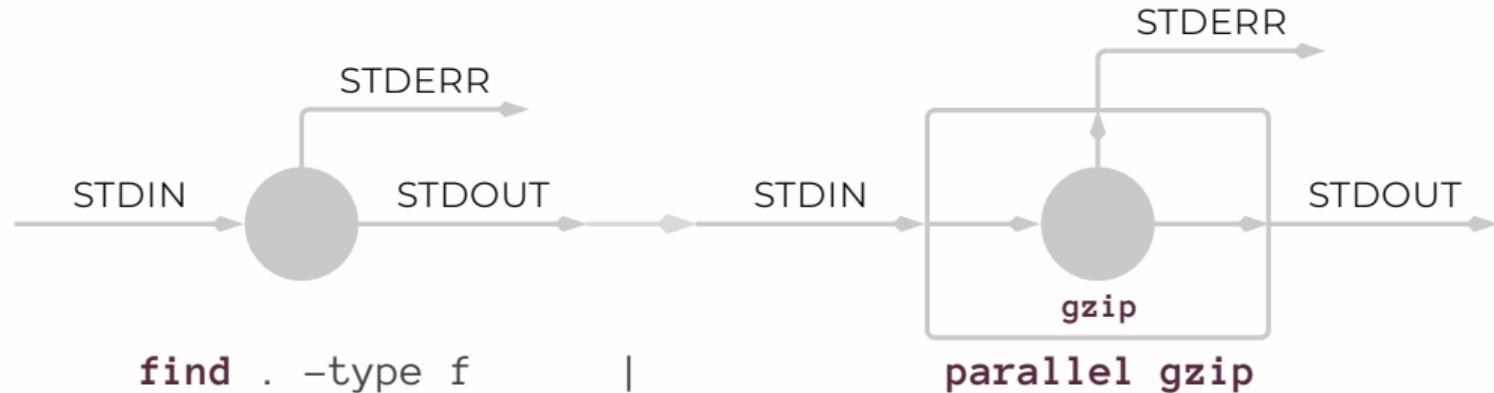
`xargs gzip`

```
./Dockerfile  
./run-docker  
./makefile  
⋮
```



```
find . -type f |
```

```
./Dockerfile  
./run-docker  
./makefile  
:
```



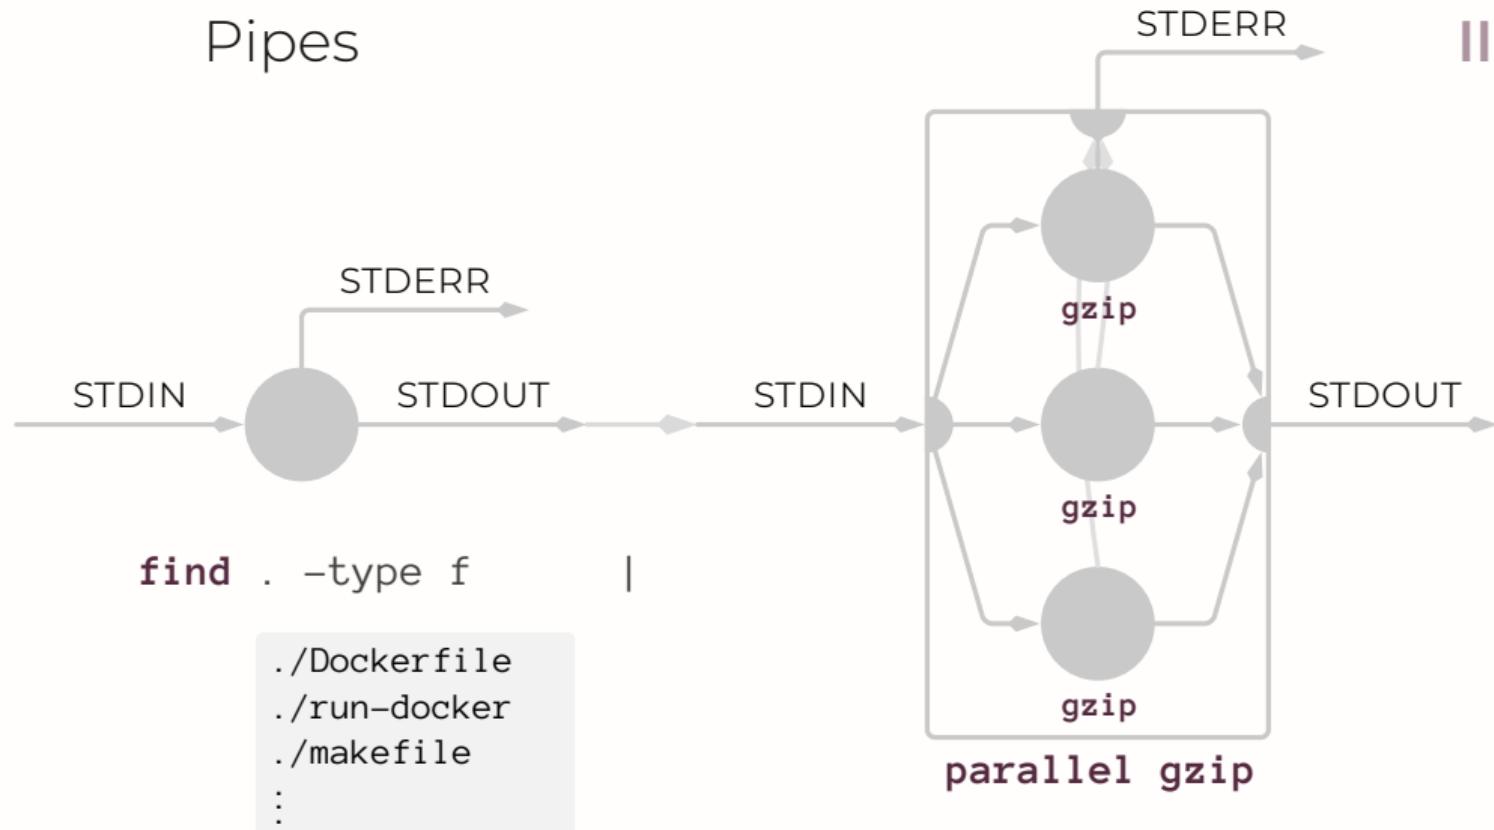
```
find . -type f | parallel gzip
```

```
./Dockerfile  
./run-docker  
./makefile  
⋮
```

5.3

## Pipes

III



**6.1**

## A simple Bank



Motivation



Background



**GNU parallel**



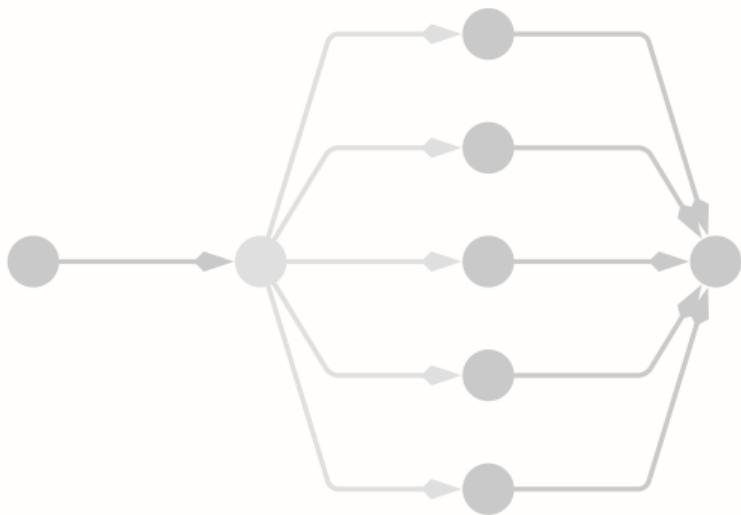
Inner Workings

Outlook



**6.2**

## A simple Bank



Motivation



Background



**GNU parallel**



Inner Workings

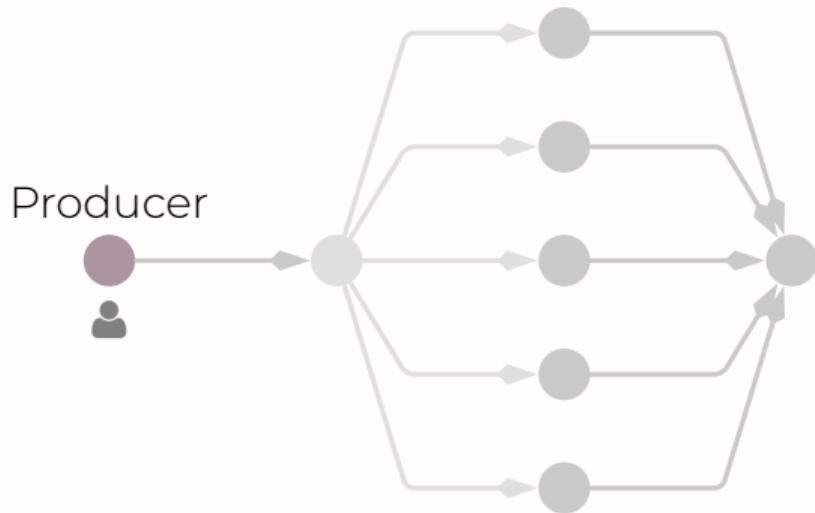


Outlook



6.3

## A simple Bank



Motivation



Background



**GNU parallel**



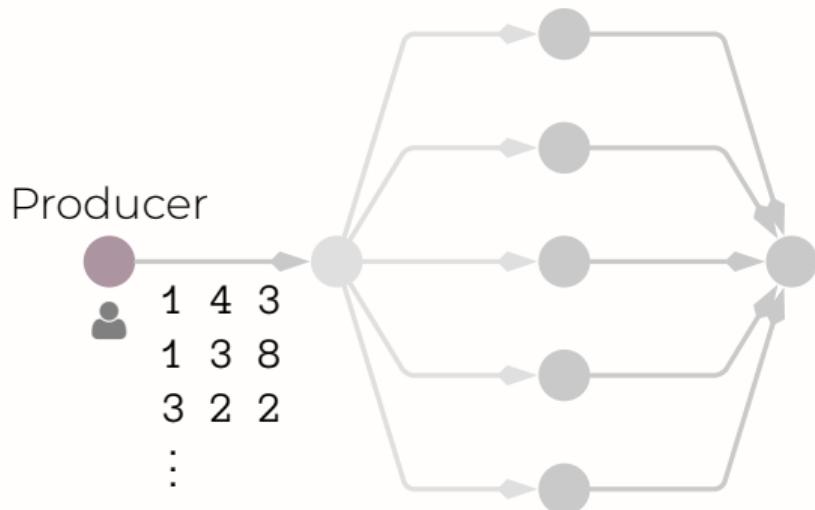
Inner Workings

Outlook



6.4

## A simple Bank



Motivation



Background



**GNU parallel**



Inner Workings

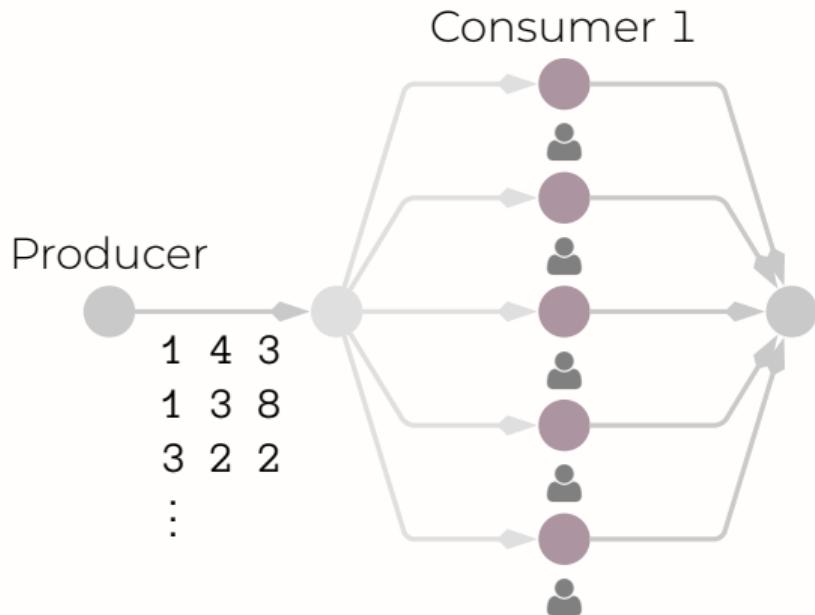


Outlook



6.5

## A simple Bank



Motivation



Background



**GNU parallel**



Inner Workings

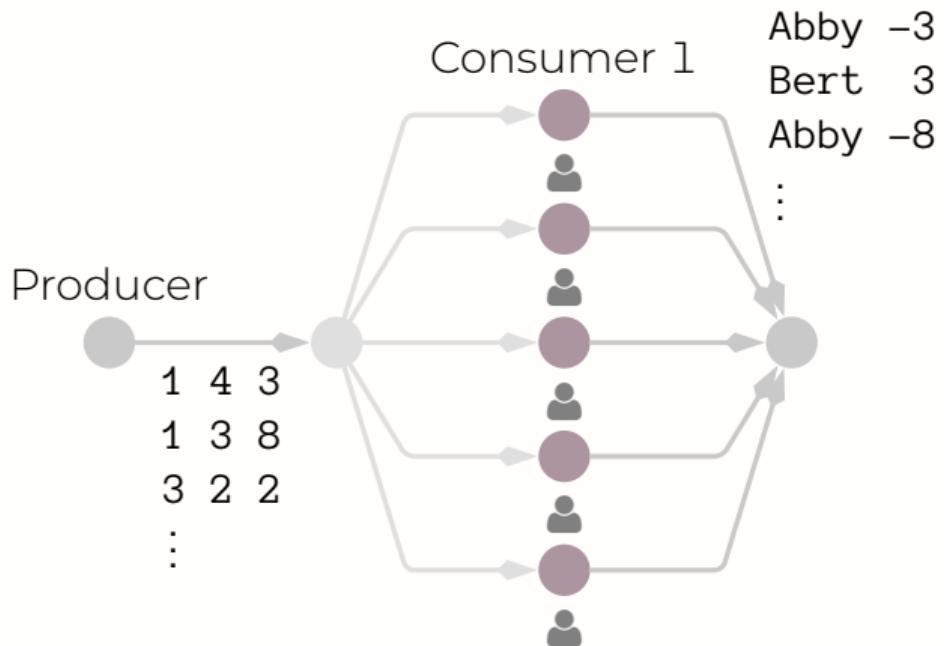


Outlook



## 6.6

# A simple Bank



Motivation



Background



**GNU parallel**



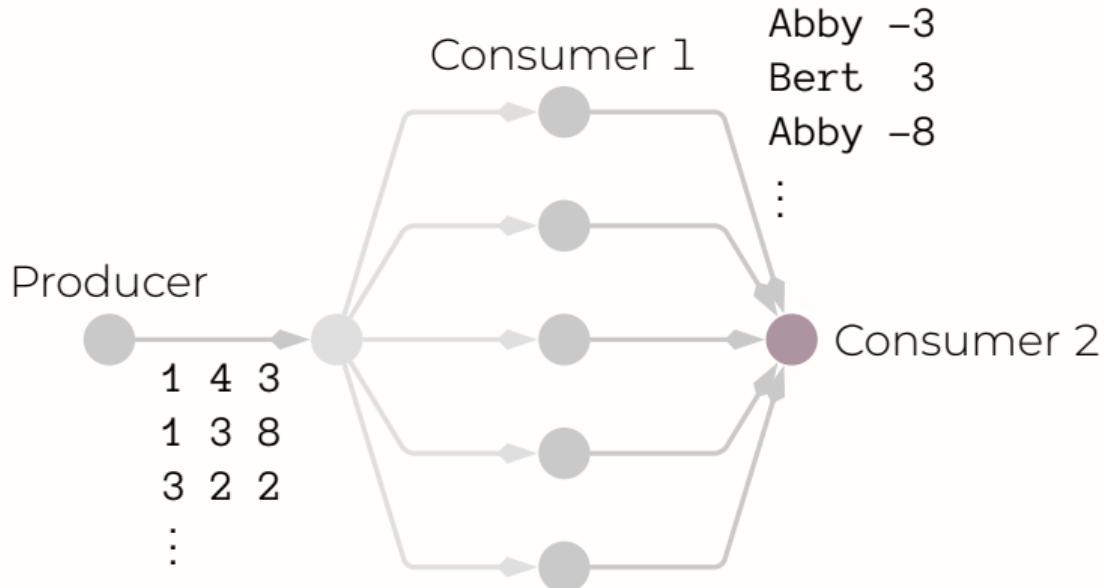
Inner Workings

Outlook



6.7

## A simple Bank



Motivation



Background



GNU parallel



Inner Workings

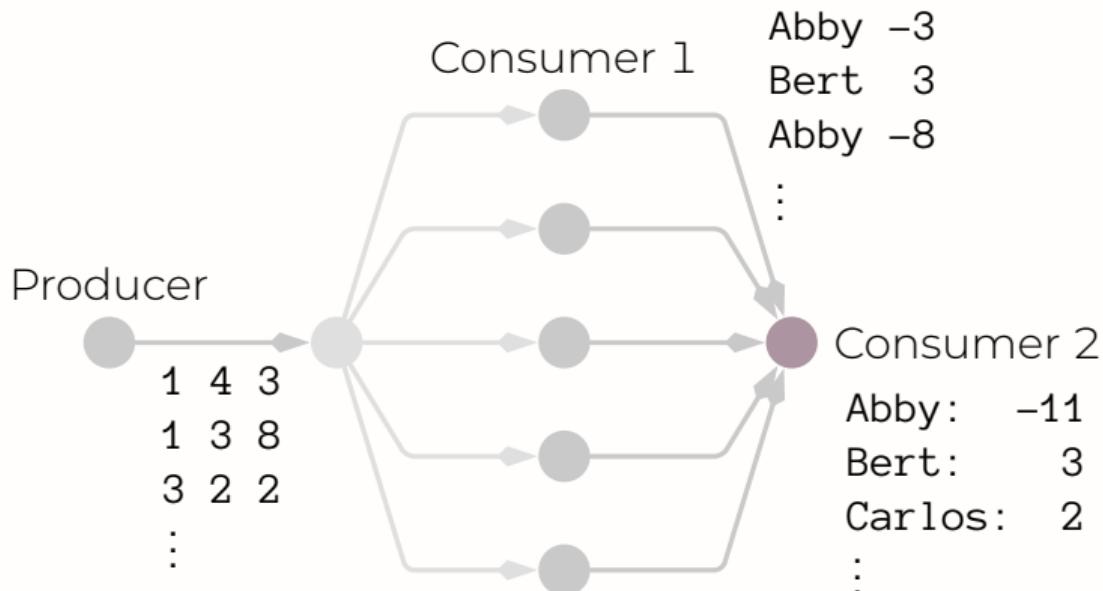


Outlook



## 6.8

# A simple Bank



Motivation



Background



GNU parallel



Inner Workings

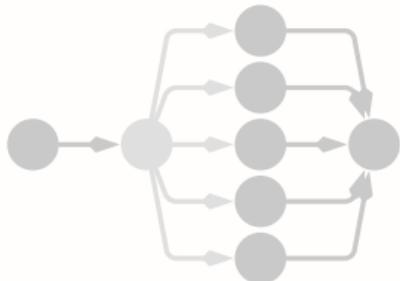


Outlook



**6.9**

## A simple Bank



Motivation



Background



**GNU parallel**



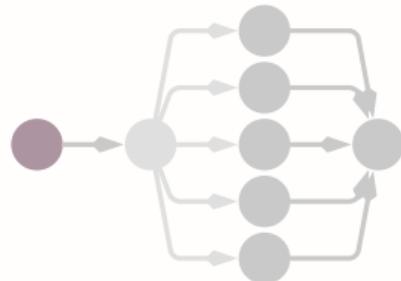
Inner Workings

Outlook



## 6.10

# A simple Bank



Producer.java

```
final var rand = new Random();

for (int i = 0; i < n; i++) {
    int from = rand.nextInt(NAMES.length);
    int to = rand.nextInt(NAMES.length);
    System.out.format("%d %d %d%n", from, to, rand.nextInt(100) * 10);
}
```



Motivation



Background



GNU parallel



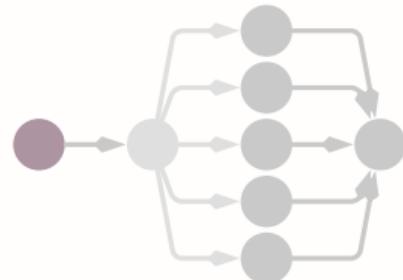
Inner Workings

Outlook



## 6.11

# A simple Bank



[Producer.java](#)

```
final var rand = new Random();
                    ↑ number of transactions
for (int i = 0; i < n; i++) {
    int from = rand.nextInt(NAMES.length);
    int to = rand.nextInt(NAMES.length);
    System.out.format("%d %d %d%n", from, to, rand.nextInt(100) * 10);
}
```



Motivation



Background



GNU parallel



Inner Workings

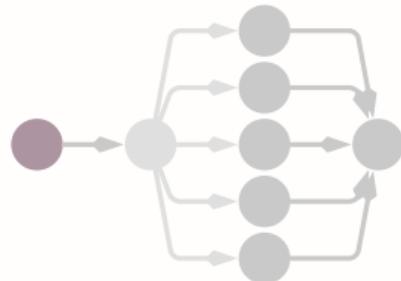


Outlook



## 6.12

# A simple Bank



[Producer.java](#)

```
final var rand = new Random();
                    ↗ number of transactions
for (int i = 0; i < n; i++) { ↗ known account names
    int from = rand.nextInt(NAMES.length);
    int to = rand.nextInt(NAMES.length);
    System.out.format("%d %d %d%n", from, to, rand.nextInt(100) * 10);
}
```



Motivation



Background



GNU parallel



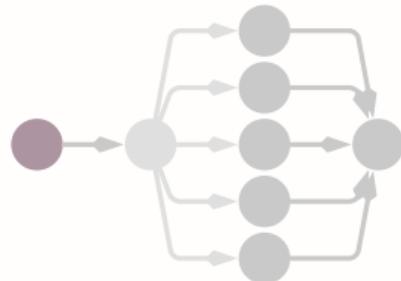
Inner Workings

Outlook



6.13

## A simple Bank



[Producer.java](#)

```
final var rand = new Random();

for (int i = 0; i < n; i++) {    ↗ number of transactions
    int from = rand.nextInt(NAMES.length);
    int to = rand.nextInt(NAMES.length);
    System.out.format("%d %d %d%n", from, to, rand.nextInt(100) * 10);
}
```

↗ known account names

value



Motivation



Background



GNU parallel



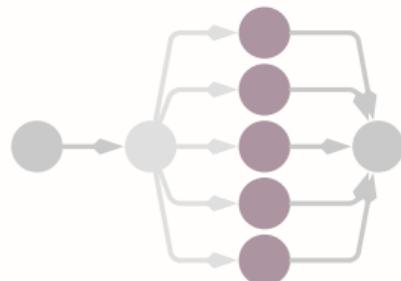
Inner Workings

Outlook



## 6.14

# A simple Bank



Consumer.java

```
final var scanner = new Scanner(System.in);
while (scanner.hasNextLine()) {
    String[] s = scanner.nextLine().split(" ");

    String from = NAMES[Integer.parseInt(s[FROM])];
    System.out.printf("%s-%d%n", from, Integer.parseInt(s[VALUE]));

    String to = NAMES[Integer.parseInt(s[TO])];
    System.out.printf("%s.%d%n", to, Integer.parseInt(s[VALUE]));
}
scanner.close();
```



Motivation



Background



GNU parallel



Inner Workings



Outlook

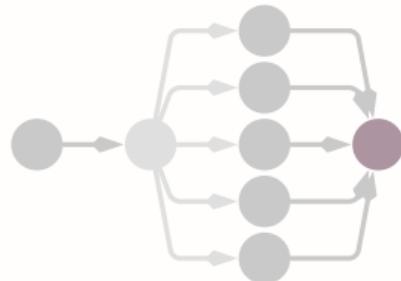


## 6.15

# A simple Bank

Accountant.java

```
final var scanner = new Scanner(System.in);
final var accounts = new HashMap<String, Integer>();
while (scanner.hasNextLine()) {
    String[] trans = scanner.nextLine().split(" ");
    final var old = accounts.getOrDefault(trans[0], 0);
    accounts.put(trans[0], old + Integer.parseInt(trans[1]));
}
scanner.close();
System.out.println(accounts);
```



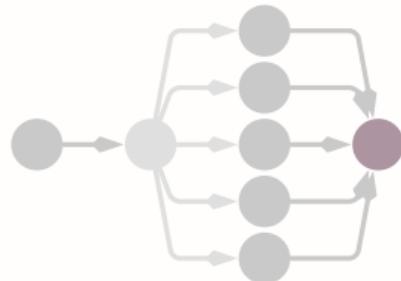
6.16

## A simple Bank

Accountant.java

```
final var scanner = new Scanner(System.in);
final var accounts = new HashMap<String, Integer>();
while (scanner.hasNextLine()) {
    String[] trans = scanner.nextLine().split(" ");
    final var old = accounts.getOrDefault(trans[0], 0);
    accounts.put(trans[0], old + Integer.parseInt(trans[1]));
}
scanner.close();
System.out.println(accounts);
```

*Arbitrary initialization*



Motivation



Background



GNU parallel



Inner Workings



Outlook



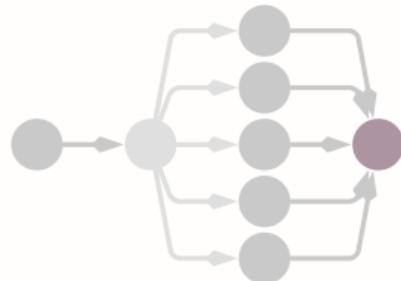
6.17

## A simple Bank

Accountant.java

```
final var scanner = new Scanner(System.in);
final var accounts = new HashMap<String, Integer>();
while (scanner.hasNextLine()) {
    String[] trans = scanner.nextLine().split(" ");
    final var old = accounts.getOrDefault(trans[0], 0);
    accounts.put(trans[0], old + Integer.parseInt(trans[1]));
}
scanner.close();
System.out.println(accounts);
```

*Arbitrary initialization*



Motivation



Background



GNU parallel



Inner Workings



Outlook



**7.1**

## Running the Example



Motivation



Background



**GNU parallel**



Inner Workings

Outlook



## 7.2

# Running the Example



Motivation



Background



**GNU parallel**



Inner Workings

Outlook



## 7.3

# Running the Example



```
java -jar producer.jar 1000000
```



Motivation

> Background

> **GNU parallel**



> Inner Workings

> Outlook



## 7.4

# Running the Example



```
java -jar producer.jar 1000000 \
| java -jar consumer.jar
```



Motivation



Background



GNU parallel



Inner Workings

Outlook



## 7.5

# Running the Example



```
java -jar producer.jar 1000000 \
| java -jar consumer.jar \
| java -jar accountant.jar
```

## 7.6

# Running the Example



```
java -jar producer.jar 1000000 \
| java -jar consumer.jar \
| java -jar accountant.jar
```

| pid   | ppid  | cpuid | cmd   |
|-------|-------|-------|---|
| 25621 | 25618 | 7     | /bin/bash -c java -jar producer.jar 1000000   java -jar consumer.jar   java -jar accountant.jar |



## 7.7

## Running the Example



```
java -jar producer.jar 1000000 \
    | java -jar consumer.jar \
    | java -jar accountant.jar
```

| pid   | ppid  | cpuid | cmd   |
|-------|-------|-------|---|
| 25621 | 25618 | 7     | /bin/bash -c java -jar producer.jar 1000000   java -jar consumer.jar   java -jar accountant.jar |
| 25622 | 25621 | 2     | java -jar producer.jar 1000000  |
| 25623 | 25621 | 14    | java -jar consumer.jar  |
| 25624 | 25621 | 13    | java -jar accountant.jar  |
| 25738 | 5113  | 9     | ps -o pid,ppid,cpuid,cmd  |



## 7.8

# Running the Example



```
java -jar producer.jar 1000000 \
    | java -jar consumer.jar \
    | java -jar accountant.jar
```

| pid   | ppid  | cpuid | cmd   |
|-------|-------|-------|---|
| 25621 | 25618 | 7     | /bin/bash -c java -jar producer.jar 1000000   java -jar consumer.jar   java -jar accountant.jar |
| 25622 | 25621 | 2     | java -jar producer.jar 1000000  |
| 25623 | 25621 | 14    | java -jar consumer.jar  |
| 25624 | 25621 | 13    | java -jar accountant.jar  |
| 25738 | 5113  | 9     | ps -o pid,ppid,cpuid,cmd  |



## 7.9

# Running the Example



```
java -jar producer.jar 1000000\  
| java -jar consumer.jar  
| java -jar accountant.jar
```

| pid   | ppid  | cpuid | cmd   |
|-------|-------|-------|---|
| 25621 | 25618 | 7     | /bin/bash -c java -jar producer.jar 1000000   java -jar consumer.jar   java -jar accountant.jar |
| 25622 | 25621 | 2     | <b>java -jar producer.jar 1000000</b>   |
| 25623 | 25621 | 14    | <b>java -jar consumer.jar</b>   |
| 25624 | 25621 | 13    | <b>java -jar accountant.jar</b>   |
| 25738 | 5113  | 9     | ps -o pid,ppid,cpuid,cmd  |



Motivation



Background

**GNU parallel**

Inner Workings

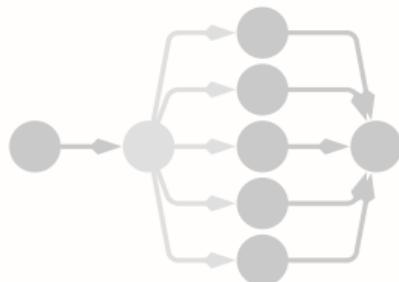


Outlook



## 7.10

## Running the Example



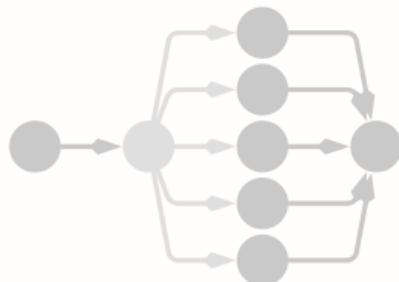
```
java -jar producer.jar 1000000\  
| parallel --pipe java -jar consumer.jar\  
| java -jar accountant.jar
```

| pid   | ppid  | cpuid | cmd   |
|-------|-------|-------|---|
| 25621 | 25618 | 7     | /bin/bash -c java -jar producer.jar 1000000   java -jar consumer.jar   java -jar accountant.jar |
| 25622 | 25621 | 2     | <b>java -jar producer.jar 1000000</b>   |
| 25623 | 25621 | 14    | <b>java -jar consumer.jar</b>   |
| 25624 | 25621 | 13    | <b>java -jar accountant.jar</b>   |
| 25738 | 5113  | 9     | ps -o pid,ppid,cpuid,cmd  |



## 7.11

# Running the Example



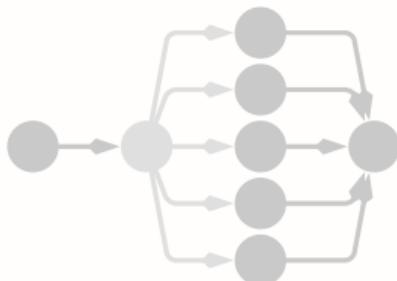
```
java -jar producer.jar 1000000\  
| parallel --pipe java -jar consumer.jar\  
| java -jar accountant.jar
```

| pid   | ppid  | cpuid | cmd   |
|-------|-------|-------|---|
| 38279 | 38276 | 3     | /bin/bash -c java -jar producer.jar 1000000   parallel --pipe java -jar consumer.jar   java -jar a... |
| ...   | ...   | ...   | ...   |



7.12

## Running the Example



```
java -jar producer.jar 1000000\  
| parallel --pipe java -jar consumer.jar\  
| java -jar accountant.jar
```

| pid   | ppid  | cpuid | cmd  |
|-------|-------|-------|--|
| 38279 | 38276 | 3     | /bin/bash -c java -jar producer.jar 1000000   <b>parallel</b> --pipe java -jar consumer.jar   java -jar a... |
| 38280 | 38279 | 6     | java -jar producer.jar 1000000   |
| 38281 | 38279 | 15    | perl /usr/bin/parallel --pipe java -jar consumer.jar   |
| 38282 | 38279 | 5     | java -jar accountant.jar   |
| 38344 | 38281 | 10    | perl -e if(sysread(STDIN,\$buf,1)){open(\$fh [...])} /usr/bin/bash -c java -jar consumer.jar                 |
| 38345 | 38281 | 4     | perl -e if(sysread(STDIN,\$buf,1)){open(\$fh [...])} /usr/bin/bash -c java -jar consumer.jar                 |
| ...   | ...   | ...   | ...  |
| 38363 | 38281 | 10    | perl -e if(sysread(STDIN,\$buf,1)){open(\$fh [...])} /usr/bin/bash -c java -jar consumer.jar                 |
| 38383 | 38281 | 4     | perl /usr/bin/parallel --pipe java -jar consumer.jar   |
| 38384 | 38344 | 2     | java -jar consumer.jar   |
| 38438 | 38136 | 12    | ps -o pid,ppid,cpuid,cmd   |



Motivation

&gt; Background

> **GNU parallel**

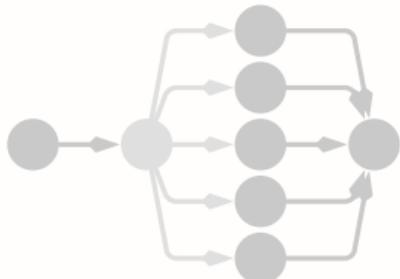
&gt; Inner Workings

&gt; Outlook



7.13

## Running the Example



```
java -jar producer.jar 1000000\  
| parallel --pipe java -jar consumer.jar\  
| java -jar accountant.jar
```

| pid   | ppid  | cpuid | cmd   |
|-------|-------|-------|---|
| 38279 | 38276 | 3     | /bin/bash -c java -jar producer.jar 1000000   parallel --pipe java -jar consumer.jar   java -jar a... |
| 38280 | 38279 | 6     | java -jar producer.jar 1000000  |
| 38281 | 38279 | 15    | perl /usr/bin/parallel --pipe java -jar consumer.jar  |
| 38282 | 38279 | 5     | java -jar accountant.jar  |
| 38344 | 38281 | 10    | perl -e if(sysread(STDIN,\$buf,1)){open(\$fh [...])} /usr/bin/bash -c java -jar consumer.jar          |
| 38345 | 38281 | 4     | perl -e if(sysread(STDIN,\$buf,1)){open(\$fh [...])} /usr/bin/bash -c java -jar consumer.jar          |
| ...   | ...   | ...   | ...   |
| 38363 | 38281 | 10    | perl -e if(sysread(STDIN,\$buf,1)){open(\$fh [...])} /usr/bin/bash -c java -jar consumer.jar          |
| 38383 | 38281 | 4     | perl /usr/bin/parallel --pipe java -jar consumer.jar  |
| 38384 | 38344 | 2     | java -jar consumer.jar  |
| 38438 | 38136 | 12    | ps -o pid,ppid,cpuid,cmd  |



Motivation

&gt; Background

> **GNU parallel**

&gt; Inner Workings

&gt; Outlook



**9.1**

# Distributed

## 9.2

# Distributed

- GNU parallel can run jobs on remote servers

## Distributed

- GNU parallel can run jobs on remote servers
  - It uses ssh to communicate with the remote machines

## Distributed

- GNU parallel can run jobs on remote servers
    - It uses ssh to communicate with the remote machines
- ```
parallel -S $SERVER echo running on :::: $SERVER
```

## Distributed

- GNU parallel can run jobs on remote servers
  - It uses ssh to communicate with the remote machines

```
parallel -S $SERVER echo running on :::: $SERVER
```
- Transfer Files using rsync:

- GNU parallel can run jobs on remote servers
  - It uses ssh to communicate with the remote machines

```
parallel -S $SERVER echo running on :::: $SERVER
```

- Transfer Files using rsync:
  - Long version:

```
parallel -S 1/"sshpass -p '$SECRET_PW' -<-->
           ssh://limerent@localhost" --transferfile {} \
           --return {}.gz --cleanup gzip :::: README.txt
```

## Distributed

- GNU parallel can run jobs on remote servers
  - It uses ssh to communicate with the remote machines

```
parallel -S $SERVER echo running on :::: $SERVER
```
- Transfer Files using rsync:
  - Long version:

```
parallel -S 1/"sshpass -p '$SECRET_PW' -< ssh://limerent@localhost" --transferfile {} \  
--return {}.gz --cleanup gzip :::: README.txt
```
  - Shorthands like -trc (transferfile, return, cleanup)

**11.1**

## Other Languages



## 11.2

# Other Languages

```
consumer.ts

const accMap = new Map<string, number>();
const rl = createInterface({ input: process.stdin });

function handleLine(input: string): void {
    const s = input.split(" ");
    accMap.set(s[0], accMap.get(s[0]) ?? 0 + Number(s[1]));
}

rl.on("line", handleLine);
rl.on("close", () => console.log(accMap));
```

## 11.3

# Other Languages

consumer.ts

```
const accMap = new Map<string, number>(); Read from STDIN
const rl = createInterface({ input: process.stdin });

function handleLine(input: string): void {
    const s = input.split(" ");
    accMap.set(s[0], accMap.get(s[0]) ?? 0 + Number(s[1]));
}

rl.on("line", handleLine);
rl.on("close", () => console.log(accMap));
```

## 11.4

# Other Languages

consumer.ts

```
const accMap = new Map<string, number>();           ↗ Read from STDIN
const rl = createInterface({ input: process.stdin });

function handleLine(input: string): void {
    const s = input.split(" ");
    accMap.set(s[0], accMap.get(s[0]) ?? 0 + Number(s[1]));
}

rl.on("line", handleLine);
rl.on("close", () => console.log(accMap));
```

Use 0 as initial value

## 11.5

# Other Languages

consumer.ts

```
const accMap = new Map<string, number>();           ↗ Read from STDIN
const rl = createInterface({ input: process.stdin });

function handleLine(input: string): void {
    const s = input.split(" ");
    accMap.set(s[0], accMap.get(s[0]) ?? 0 + Number(s[1]));
}

rl.on("line", handleLine);
rl.on("close", () => console.log(accMap));
```

Use 0 as initial value

Register callback for 'line' event

## 11.6

# Other Languages

consumer.ts

```
const accMap = new Map<string, number>();           ↗ Read from STDIN
const rl = createInterface({ input: process.stdin });

function handleLine(input: string): void {
    const s = input.split(" ");
    accMap.set(s[0], accMap.get(s[0]) ?? 0 + Number(s[1]));
}   ↗ Use 0 as initial value

rl.on("line", handleLine);                           ↗ Register callback for 'line' event
rl.on("close", () => console.log(accMap));          ↗ Anonymous callback for 'closed' event
  (arrow style)
```

**12.1**

## Integrating TypeScript



Motivation



Background



**GNU parallel**



Inner Workings

Outlook



```
java -jar producer.jar 1000000 \
| parallel --pipe --roundRobin -j4 java -jar consumer.jar \
| yarn --silent start
```

```
java -jar producer.jar 1000000\  
| parallel --pipe --roundRobin -j4 java -jar consumer.jar \  
| yarn --silent start
```



## 12.4

# Integrating TypeScript

*Distributes records amongst all jobs.  
No longer guarantees order.*

```
java -jar producer.jar 1000000 \
| parallel --pipe --roundRobin -j4 java -jar consumer.jar \
| yarn --silent start
```



## 12.5

# Integrating TypeScript

```
java -jar producer.jar 1000000 \
| parallel --pipe --roundrobin -j4 java -jar consumer.jar \
| yarn --silent start
```

--recend "\n"

Distributes records amongst all jobs.

No longer guarantees order.

--jobs 4



## 12.6

# Integrating TypeScript

```
java -jar producer.jar 1000000 \
| parallel --pipe --roundrobin -j4 java -jar consumer.jar \
| yarn --silent start
```

--recend "\n"  
Distributes records amongst all jobs.  
No longer guarantees order.

--jobs 4  
Runs tsc && node consumer.js



## 12.7

## Integrating TypeScript

```
java -jar producer.jar 1000000 \
| parallel --pipe --roundrobin -j4 java -jar consumer.jar \
| yarn --silent start
```

--recend "\n"  
Distributes records amongst all jobs.  
No longer guarantees order.

--jobs 4  
Runs tsc && node consumer.js

| pid    | ppid   | cpuid | cmd                                                                                                |
|--------|--------|-------|----------------------------------------------------------------------------------------------------|
| 258484 | 258481 | 3     | /bin/bash -c java -jar producer.jar 1000000   parallel --pipe [...]   yarn --silent start-consumer |



## 12.8

## Integrating TypeScript

```
java -jar producer.jar 1000000 \
| parallel --pipe --roundrobin -j4 java -jar consumer.jar \
| yarn --silent start
```

--recend "\n"  
Distributes records amongst all jobs.  
No longer guarantees order.

--jobs 4  
Runs tsc && node consumer.js

| pid    | ppid   | cpuid | cmd                                                                                                |
|--------|--------|-------|----------------------------------------------------------------------------------------------------|
| 258484 | 258481 | 3     | /bin/bash -c java -jar producer.jar 1000000   parallel --pipe [...]   yarn --silent start-consumer |
| 258485 | 258484 | 2     | java -jar producer.jar 1000000                                                                     |
| 258486 | 258484 | 9     | /usr/bin/perl /usr/bin/parallel --pipe --roundrobin -j4 java -jar consumer.jar                     |
| 258487 | 258484 | 3     | node /home/lord-waddle/.nvm/versions/no:c:de/v16.15.1/bin/yarn -silent start-consumer              |
| 258519 | 258486 | 11    | java -jar consumer.jar                                                                             |
| 258520 | 258486 | 3     | java -jar consumer.jar                                                                             |
| 258522 | 258486 | 7     | java -jar consumer.jar                                                                             |
| 258524 | 258486 | 10    | java -jar consumer.jar                                                                             |
| 258604 | 258487 | 11    | /bin/sh -c tsc && node consumer.js                                                                 |
| 258605 | 258604 | 4     | /home/lord-waddle/.nvm/versions/node/v16.15.1/bin/node consumer.js                                 |
| 258618 | 258480 | 11    | ps -o pid,ppid,cpu_id,cmd                                                                          |

## 12.9

## Integrating TypeScript

```
java -jar producer.jar 1000000 \
| parallel --pipe --roundrobin -j4 java -jar consumer.jar \
| yarn --silent start
```

--recend "\n"  
*Distributes records amongst all jobs.*  
*No longer guarantees order.*

--jobs 4  
*Runs tsc && node consumer.js*

| pid    | ppid   | cpuid | cmd                                                                                                |
|--------|--------|-------|----------------------------------------------------------------------------------------------------|
| 258484 | 258481 | 3     | /bin/bash -c java -jar producer.jar 1000000   parallel --pipe [...]   yarn --silent start-consumer |
| 258485 | 258484 | 2     | java -jar producer.jar 1000000                                                                     |
| 258486 | 258484 | 9     | /usr/bin/perl /usr/bin/parallel --pipe --roundrobin -j4 java -jar consumer.jar                     |
| 258487 | 258484 | 3     | node /home/lord-waddle/.nvm/versions/no:c:de/v16.15.1/bin/yarn -silent start-consumer              |
| 258519 | 258486 | 11    | java -jar consumer.jar                                                                             |
| 258520 | 258486 | 3     | java -jar consumer.jar                                                                             |
| 258522 | 258486 | 7     | java -jar consumer.jar                                                                             |
| 258524 | 258486 | 10    | java -jar consumer.jar                                                                             |
| 258604 | 258487 | 11    | /bin/sh -c tsc && node consumer.js                                                                 |
| 258605 | 258604 | 4     | /home/lord-waddle/.nvm/versions/node/v16.15.1/bin/node consumer.js                                 |
| 258618 | 258480 | 11    | ps -o pid,ppid,cpu_id,cmd                                                                          |

*Path on Florian's system*

**14.1**

## Recap

 gRPC [07/01/22]

 protocol buffers [07/01/22]



Motivation



Background



**GNU parallel**



Inner Workings



Outlook



- Stream-based communication

- Stream-based communication
  - Cf. Java's functional streams

- Stream-based communication
  - Cf. Java's functional streams
  - Serialization and deserialization

- Stream-based communication
  - Cf. Java's functional streams
  - Serialization and deserialization
  - Decoupled programs (e.g., no shared memory)

- Stream-based communication
  - Cf. Java's functional streams
  - Serialization and deserialization
  - Decoupled programs (e.g., no shared memory)
- Allows distribution individual operators in the pipeline

- Stream-based communication
  - Cf. Java's functional streams
  - Serialization and deserialization
  - Decoupled programs (e.g., no shared memory)
- Allows distribution individual operators in the pipeline
- Easy combination of different languages

- Stream-based communication
  - Cf. Java's functional streams
  - Serialization and deserialization
  - Decoupled programs (e.g., no shared memory)
- Allows distribution individual operators in the pipeline
- Easy combination of different languages
  - Comparable with gRPC

- Stream-based communication
  - Cf. Java's functional streams
  - Serialization and deserialization
  - Decoupled programs (e.g., no shared memory)
- Allows distribution individual operators in the pipeline
- Easy combination of different languages
  - Comparable with gRPC
  - But: no message standardization (cf. protocol buffers)

- Stream-based communication
  - Cf. Java's functional streams
  - Serialization and deserialization
  - Decoupled programs (e.g., no shared memory)
- Allows distribution individual operators in the pipeline
- Easy combination of different languages
  - Comparable with gRPC
  - But: no message standardization (cf. protocol buffers)
- Programs don't know anything of the parallelization

 gRPC [07/01/22]

 protocol buffers [07/01/22]

**15.1**

# About Pipelines

**bash**

[↳ posix/pipe \[07/01/22\]](#)

[↳ unix pipeline \[07/01/22\]](#)

[↳ bash.pipelines \[07/01/22\]](#)

[↳ named pipes \[07/01/22\]](#)



Motivation



Background



GNU parallel



**Inner Workings**



Outlook



- Executes each program in own subshell

[↳ posix/pipe \[07/01/22\]](#)

[↳ unix pipeline \[07/01/22\]](#)

[↳ bash.pipelines \[07/01/22\]](#)

[↳ named pipes \[07/01/22\]](#)



Motivation



Background



GNU parallel



**Inner Workings**



Outlook



- Executes each program in own subshell
- Buffer provided by the kernel

[↳ posix/pipe \[07/01/22\]](#)

[↳ unix pipeline \[07/01/22\]](#)

[↳ bash.pipelines \[07/01/22\]](#)

[↳ named pipes \[07/01/22\]](#)

- Executes each program in own subshell
- Buffer provided by the kernel
  - Works on bytes (no known boundaries except max-size)

[↳ posix/pipe \[07/01/22\]](#)

[↳ unix pipeline \[07/01/22\]](#)

[↳ bash.pipelines \[07/01/22\]](#)

[↳ named pipes \[07/01/22\]](#)

- Executes each program in own subshell
- Buffer provided by the kernel
  - Works on bytes (no known boundaries except max-size)
  - Limited capacity (`/proc/sys/fs/pipe-max-size`)

[↳ posix/pipe \[07/01/22\]](#)

[↳ unix pipeline \[07/01/22\]](#)

[↳ bash.pipelines \[07/01/22\]](#)

[↳ named pipes \[07/01/22\]](#)

- Executes each program in own subshell
- Buffer provided by the kernel
  - Works on bytes (no known boundaries except max-size)
  - Limited capacity (`/proc/sys/fs/pipe-max-size`)
  - By default blocking read and write

[↳ posix/pipe \[07/01/22\]](#)

[↳ unix pipeline \[07/01/22\]](#)

[↳ bash.pipelines \[07/01/22\]](#)

[↳ named pipes \[07/01/22\]](#)

- Executes each program in own subshell
- Buffer provided by the kernel
  - Works on bytes (no known boundaries except max-size)
  - Limited capacity (`/proc/sys/fs/pipe-max-size`)
  - By default blocking read and write
  - Can be changed with `O_NONBLOCK` flag (`pipe2, fnctl`)

 [posix/pipe \[07/01/22\]](#)

 [unix pipeline \[07/01/22\]](#)

 [bash.pipelines \[07/01/22\]](#)

 [named pipes \[07/01/22\]](#)

- Executes each program in own subshell
- Buffer provided by the kernel
  - Works on bytes (no known boundaries except max-size)
  - Limited capacity (`/proc/sys/fs/pipe-max-size`)
  - By default blocking read and write
  - Can be changed with `O_NONBLOCK` flag (`pipe2, fcntl`)
  - This sets `errno` to `EWOULDBLOCK` or `EAGAIN`

 [posix/pipe \[07/01/22\]](#)

 [unix pipeline \[07/01/22\]](#)

 [bash.pipelines \[07/01/22\]](#)

 [named pipes \[07/01/22\]](#)

- Executes each program in own subshell
- Buffer provided by the kernel
  - Works on bytes (no known boundaries except max-size)
  - Limited capacity (`/proc/sys/fs/pipe-max-size`)
  - By default blocking read and write
  - Can be changed with `O_NONBLOCK` flag (`pipe2, fnctl`)
  - This sets `errno` to `EWOULDBLOCK` or `EAGAIN`
- By default unidirectional

[↳ posix/pipe \[07/01/22\]](#)

[↳ unix pipeline \[07/01/22\]](#)

[↳ bash.pipelines \[07/01/22\]](#)

[↳ named pipes \[07/01/22\]](#)

- Executes each program in own subshell
- Buffer provided by the kernel
  - Works on bytes (no known boundaries except max-size)
  - Limited capacity (`/proc/sys/fs/pipe-max-size`)
  - By default blocking read and write
  - Can be changed with `O_NONBLOCK` flag (`pipe2, fnctl`)
  - This sets `errno` to `EWOULDBLOCK` or `EAGAIN`
- By default unidirectional
  - Named pipes (like `fifo`) allow half duplex data flow

[↳ posix/pipe \[07/01/22\]](#)

[↳ unix pipeline \[07/01/22\]](#)

[↳ bash.pipelines \[07/01/22\]](#)

[↳ named pipes \[07/01/22\]](#)

**16.1**

## GNU Parallel

 parallel\_design [07/01/22]

 unbuffered output [07/01/22]



Motivation



Background



GNU parallel



**Inner Workings**

> Outlook



- Mixes tabs and spaces for padding 

- Mixes tabs and spaces for padding 
- Supplied as a single file (object-oriented Perl)

- Mixes tabs and spaces for padding 
- Supplied as a single file (object-oriented Perl)
  - Runs wherever there is a Perl interpreter

- Mixes tabs and spaces for padding 
- Supplied as a single file (object-oriented Perl)
  - Runs wherever there is a Perl interpreter
  - Rather slow, 3–10 ms per job and 1 ms/MB output

- Mixes tabs and spaces for padding 
- Supplied as a single file (object-oriented Perl)
  - Runs wherever there is a Perl interpreter
  - Rather slow, 3–10 ms per job and 1 ms/MB output
  - Uses busy wait (with exponential sleeping times)

- Mixes tabs and spaces for padding 
- Supplied as a single file (object-oriented Perl)
  - Runs wherever there is a Perl interpreter
  - Rather slow, 3–10 ms per job and 1 ms/MB output
  - Uses busy wait (with exponential sleeping times)
  - A lot of support for the hosting shell

- Mixes tabs and spaces for padding 
- Supplied as a single file (object-oriented Perl)
  - Runs wherever there is a Perl interpreter
  - Rather slow, 3–10 ms per job and 1 ms/MB output
  - Uses busy wait (with exponential sleeping times)
  - A lot of support for the hosting shell
- Buffers output on disk for distinction

- Mixes tabs and spaces for padding 
- Supplied as a single file (object-oriented Perl)
  - Runs wherever there is a Perl interpreter
  - Rather slow, 3–10 ms per job and 1 ms/MB output
  - Uses busy wait (with exponential sleeping times)
  - A lot of support for the hosting shell
- Buffers output on disk for distinction
- GNU parallel parses processes everything from stdin

- Mixes tabs and spaces for padding 
- Supplied as a single file (object-oriented Perl)
  - Runs wherever there is a Perl interpreter
  - Rather slow, 3–10 ms per job and 1 ms/MB output
  - Uses busy wait (with exponential sleeping times)
  - A lot of support for the hosting shell
- Buffers output on disk for distinction
- GNU parallel parses processes everything from stdin
  - E.g., this limits the throughput of --pipe

 [parallel\\_design \[07/01/22\]](#)

 [unbuffered output \[07/01/22\]](#)

**17.1**

It can do more!

 [parallel tutorial \[06/30/22\]](#)

 [parallel alternatives \[06/30/22\]](#)



Motivation



Background



GNU parallel



Inner Workings



**Outlook**



- Different spreading strategies (`--shard`, `--bin`, ...) for `--pipe`

- Different spreading strategies (--shard, --bin, ...) for --pipe
- Replacement strings ({ }, { % }, ...)

## 17.4

# It can do more!

- Different spreading strategies (`--shard`, `--bin`, ...) for `--pipe`
- Replacement strings (`{ }`, `{%}`, ...)
- Compression of buffer data (`--compress`)

- Different spreading strategies (--shard, --bin, ...) for --pipe
- Replacement strings ({ }, { % }, ...)
- Compression of buffer data (--compress)
- Comfort-Support for named pipes (--fifo)

- Different spreading strategies (--shard, --bin, ...) for --pipe
- Replacement strings ({ }, { % }, ...)
- Compression of buffer data (--compress)
- Comfort-Support for named pipes (--fifo)
- Support for unfair counting semaphore with timeout

- Different spreading strategies (`--shard`, `--bin`, ...) for `--pipe`
- Replacement strings (`{ }`, `{%}`, ...)
- Compression of buffer data (`--compress`)
- Comfort-Support for named pipes (`--fifo`)
- Support for unfair counting semaphore with timeout
  - With options (`--semaphore`)

- Different spreading strategies (--shard, --bin, ...) for --pipe
- Replacement strings ({ }, { % }, ...)
- Compression of buffer data (--compress)
- Comfort-Support for named pipes (--fifo)
- Support for unfair counting semaphore with timeout
  - With options (--semaphore)
  - As alternative Program (sem)

- Different spreading strategies (`--shard`, `--bin`, ... ) for `--pipe`
- Replacement strings (`{ }`, `{%}`, ...)
- Compression of buffer data (`--compress`)
- Comfort-Support for named pipes (`--fifo`)
- Support for unfair counting semaphore with timeout
  - With options (`--semaphore`)
  - As alternative Program (`sem`)
- Load Balancing (`--limit`, `--load`, ...)

- Different spreading strategies (`--shard`, `--bin`, ... ) for `--pipe`
- Replacement strings (`{ }`, `{%}`, ...)
- Compression of buffer data (`--compress`)
- Comfort-Support for named pipes (`--fifo`)
- Support for unfair counting semaphore with timeout
  - With options (`--semaphore`)
  - As alternative Program (`sem`)
- Load Balancing (`--limit`, `--load`, ...)
- And so much more (Tables, SQL, Shebang, ...)

[1] Ole Tange. *GNU Parallel 20210822 ('Kabul')*. 2021

**Felix R. & Florian S.**  
Ulm July 3, 2022

