



# SEA JUBE Benchmarking Workshop

DR. MAX HOLICKI | JÜLICH SUPERCOMPUTING CENTRE

# THE PEOPLE BEHIND JSC SEA BENCHMARKING?

- Sebastian Lührs (Developer of JUBE)
- Andreas Smolenko (DEEP-SEA T1.2 Lead, IO-SEA)
- Max Holicki (DEEP-SEA, IO-SEA T1.2 Lead, Red-SEA)
- Jan-Oliver Mirus (DEEP-SEA, IO-SEA)
- Yannik Müller (DEEP-, IO- & RED-SEA)
  
- Tom Ridley
- Filipe Guimarães

# SCHEDULE

08:50 Video Conference Opens

09:00 Introduction

09:15 Introduction to JUBE by Sebastian Lührs

10:15 Workshop Part 1 (installation + Hands on web tutorial)

12:00 Lunch

13:00 SEA Benchmarking + Discussion

14:00 Workshop Part 2 (application integration)

16:00 Wrap-Up

16:30 Video Conference Closes

# WHAT IS A BENCHMARK?

Origin from surveying (1884): Bench + Mark

1. Bench: An angle iron (used to support a leveling staff)
2. Mark: A marking identifying a location

Combined these define a point of reference or a point of comparison.

Nowadays benchmarking no longer refers to the process of establishing a benchmark, but to the comparing to a benchmark.



<https://collection.sciencemuseumgroup.org.uk/object/s/co52873/angle-iron-to-support-levelling-staff-at-bench-mar-bracket-structural-element-bench-mark>

, CC BY-SA 4.0,  
[org.uk/index.php?curid=57](https://www.sciencemuseumgroup.org.uk/index.php?curid=57)



# WHAT IS A BENCHMARK IN THE COMPUTER SCIENCES?

Benchmarking in the Computer Sciences is the act of comparing the execution of a program based on metrics, like runtime.

Often this is done to ensure that there are no regressions in performance, which makes it akin to regression testing, but not 100% the same.

# WHAT DO WE WANT TO ACHIEVE IN IO-SEA?

# Improve IO

# WHAT DOES THIS MEAN FOR BENCHMARKING IN IO-SEA?

Network  
Bandwidth

IO-Time

IO-Latency


## Benchmark IO

IO-Bandwidth

Runtime

Network  
Latency

# WHAT IS OUR END GOAL?



This graph is purely illustrative!



# WHY BENCHMARK IN SEA PROJECTS?

1. **Accountability:** At the end of the day we need to demonstrate to our sponsors that the money invested into us has led to something and that this is also an agreed possible outcome. *Trust is good, control is better!*
2. **Staying-on-Track/Regression-Testing:** Consistent benchmarking will allow us to stay on track and give feedback to developers if their codes/systems are improving. *In a sense we do regression testing for you.*

# JUBE TUTORIAL

- Online at: <https://apps.fz-juelich.de/jsc/jube/jube2/docu/tutorial.html>
- On JSC systems:
  1. Load JUBE: *module load JUBE/2.4.1*
  2. Copy Examples: *cp \$EBROOTJUBE/examples <PATH>*
- From TAR file:
  1. Decompress archive
  2. Navigate to examples folder
- Reservation: workshop\_hpc on JUSUF



# SEA JUBE Benchmarking

DR. MAX HOLICKI | JÜLICH SUPERCOMPUTING CENTRE

# THE PLAN I

- Benchmark systems using use cases and synthetic benchmarks.
- Why synthetic benchmarks?
  - We need to understand performance change in use cases.
    - Are they due to use-case–code changes?
    - Are they due to software changes?
    - Are they due to hardware changes?

# THE PLAN II

- Weekly or Biweekly Schedule.
- Automated Benchmarking via bots, cron jobs and JUBE.
- Aggregate results on GitLab.
  - One repository per software, one branch for each system.
- Archive important benchmark results on local machines.
  
- **We need you to keep a log of system/software changes!**

# THE CRON JOB

## (Bi-)Weekly Schedule

### JUBE

1. Build
2. Benchmark
3. Extract Results



### GitLab

1. Checkout Branch
2. Push Results



### Tar

1. Compress Benchmark Results



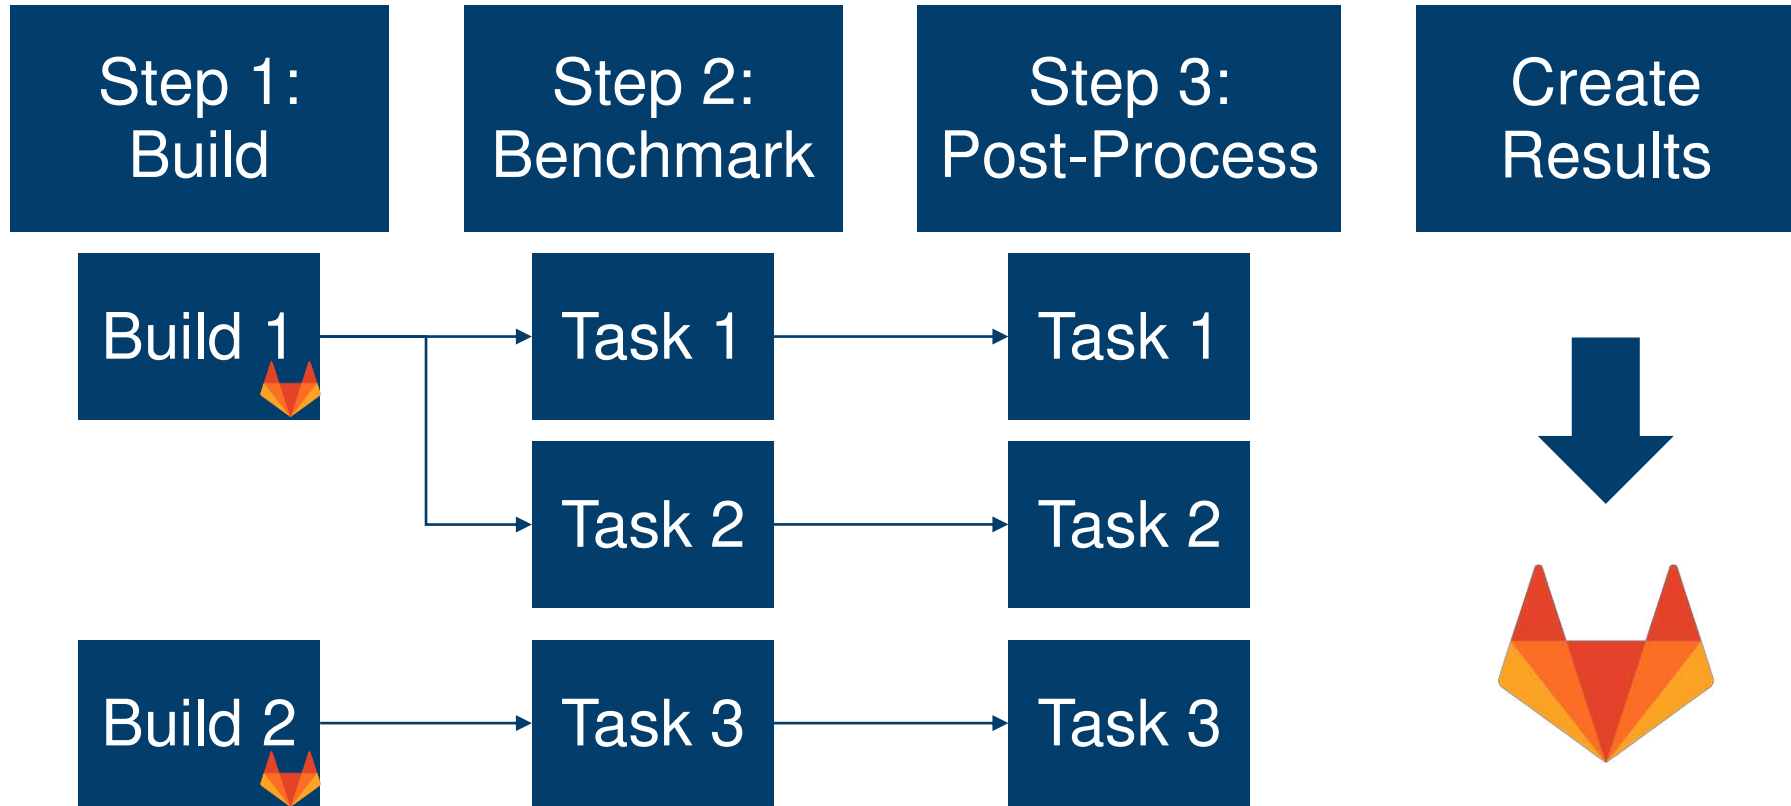
Von GitLab B.V. - [gitlab.com/gitlab-org/gitlab-ce/blob/master/app/assets/images/logo\\_wordmark.svg](https://gitlab.com/gitlab-org/gitlab-ce/blob/master/app/assets/images/logo_wordmark.svg), MIT, <https://commons.wikimedia.org/w/index.php?curid=74127414>

Von rar-labs - <http://www.seeklogo.com/files/W/WinRAR-vector-logo-9A6D30BEBE-seeklogo.com.zip> /  
farben von Datei:WinRAR-Logo.png, Logo, <https://de.wikipedia.org/w/index.php?curid=6203168>

# THE BENCHMARKS

- Each Benchmark will consist of at least two steps:
  - Step 1: Build (compile, download files & executable, etc.).  
We always want to be using the freshest build possible.
  - Step 2: Run benchmark.
  - Step 3: (Optional) post-process results.
- After benchmark completion result tables are generated.
  - 1 master table is appended to.
  - 1 separate table named after the day is also generated.
- These tables are then pushed to GitLab.

# THE BENCHMARKS





# GITLAB

- One Group per project (DEEP- and IO-SEA groups present).
- One repository per use-case software or synthetic benchmark.
  - One branch per system (not an issue for DEEP).
  - Access to use-case repositories restricted to benchmarking team and use-case developers.
  - Everyone has access to synthetic benchmarks, master branch edits restricted to benchmarking team.

# BOTS

- Benchmarks will be executed on a schedule by bots.
  - Weekly or bi-weekly (TBD)
- Result post processing by the bot is possible.
- Aggregated results will automatically be uploaded to GitLab.
- All pertinent benchmark data is stored locally on the system.
  - These data can be archived. They just need to be accessible for inspection at a later date.
  - The data can be deleted at project end.

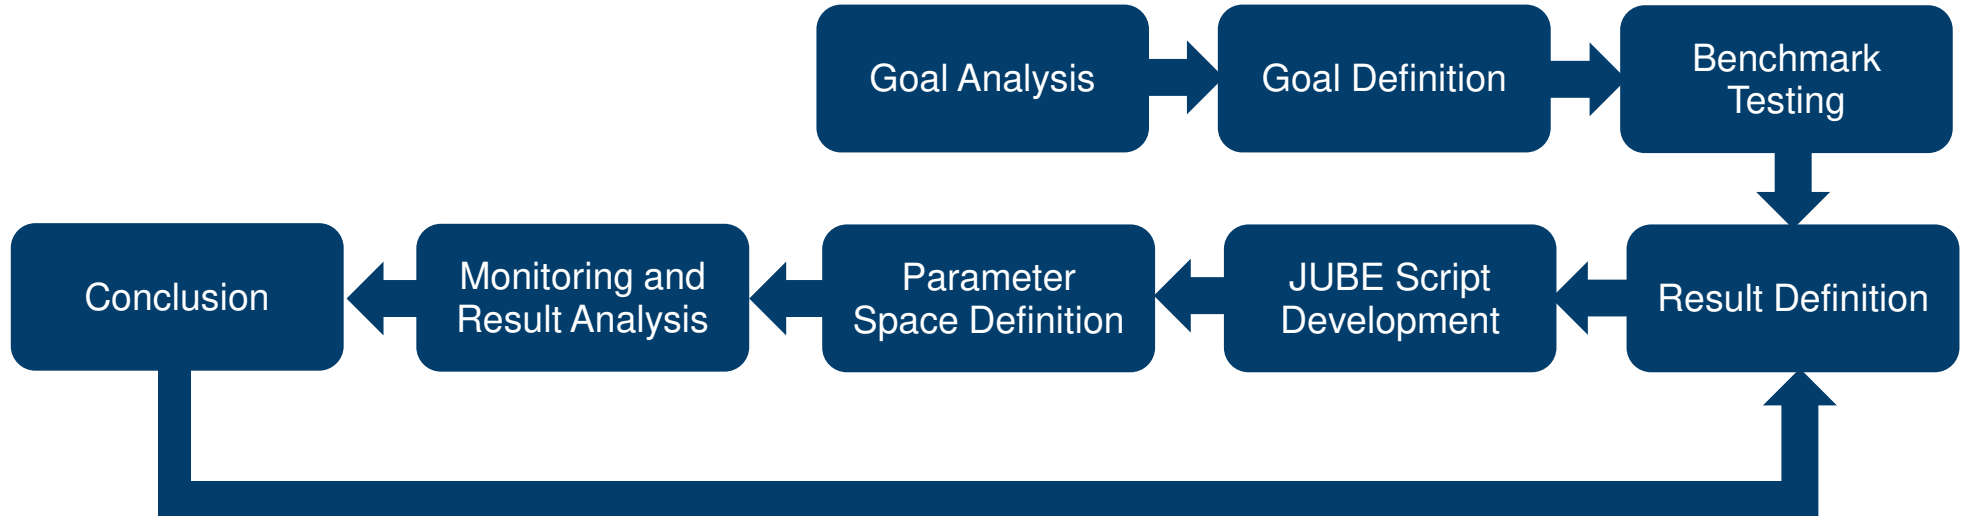
# BOTS & GITLAB

- The Bots will have GitLab accounts.
- This is mostly so that we can distinguish between user-generated benchmark data, which you are more than welcome to generate, and those created on a schedule.
- GitLab access will be provided via tokens.

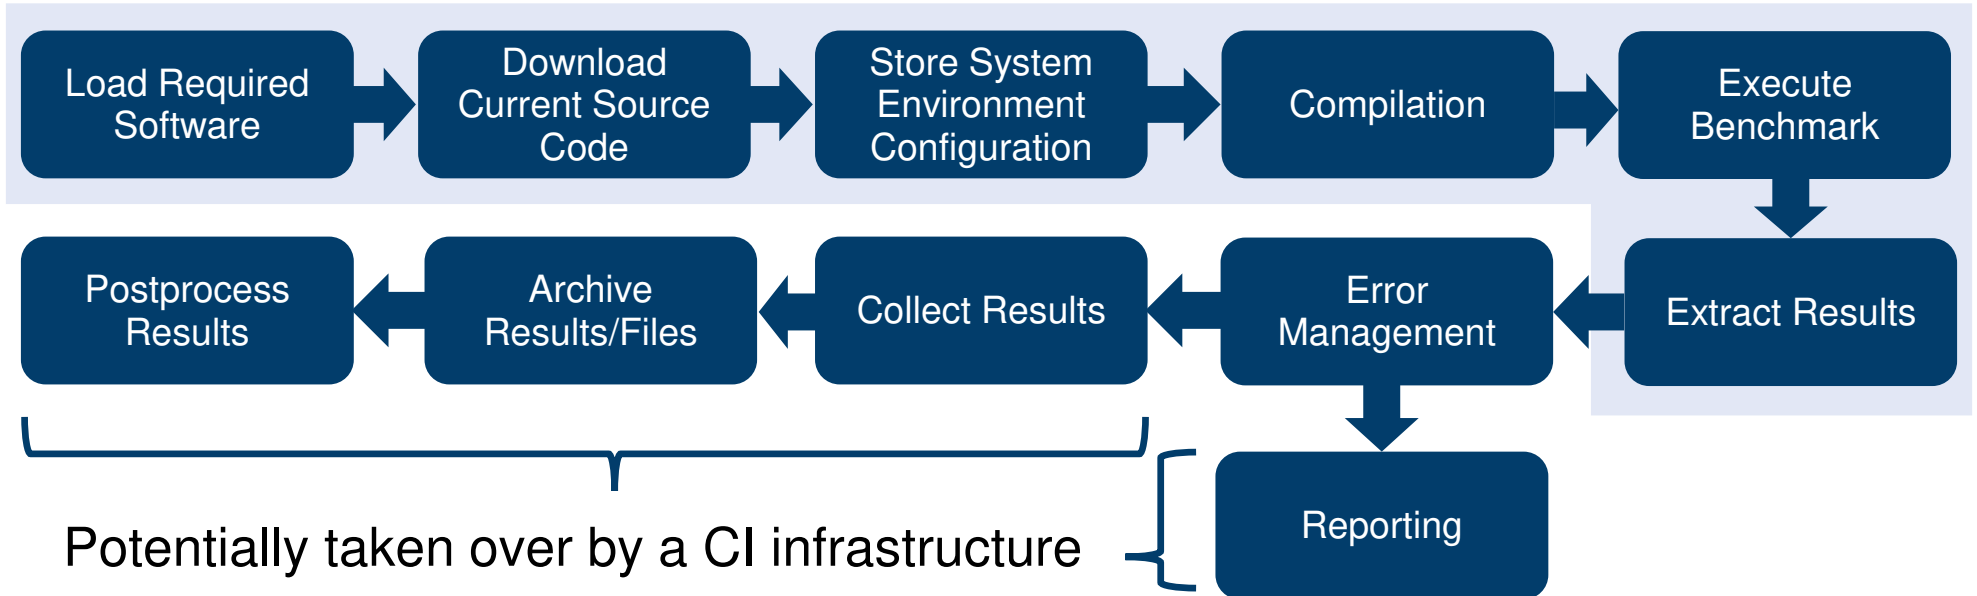
```
# 0. Parameters
username  ="deep_bench";           #The bot that will be reporting the benchmark results.
usetoken  ="????????????????????"; #The access token of the bot.
usermail  ="sc@fz-juelich.de";     #The bot does not have an email address so we use JSC support.
repository="gitlab.version.fz-juelich.de/holickil/benchmarking-workshop-linktest.git";

# 1. Switch to deep_bench user for reporting
git config user.name  ${username}; #Switch user
git config user.email ${usermail}; #Switch email
# 2. Change remote url to use access token
git remote set-url origin https://${username}:${usetoken}@${repository};
# 3. Create system repository
git checkout -b ${SYSTEMNAME};
```

# SAMPLE BENCHMARK INTEGRATION WORKFLOW





# SAMPLE BENCHMARK EXECUTION WORKFLOW



# MAKING EFFECTIVE BENCHMARKS

- KISS (Keep It Simple Stupid)
- Only add as much complexity as necessary
- Benchmarks should be short and to the point
- Time things as precisely as possible
- Output intelligently to log files
- Make benchmark quantities easily findable
- Implement proper error checking at the end
- Have Fun!

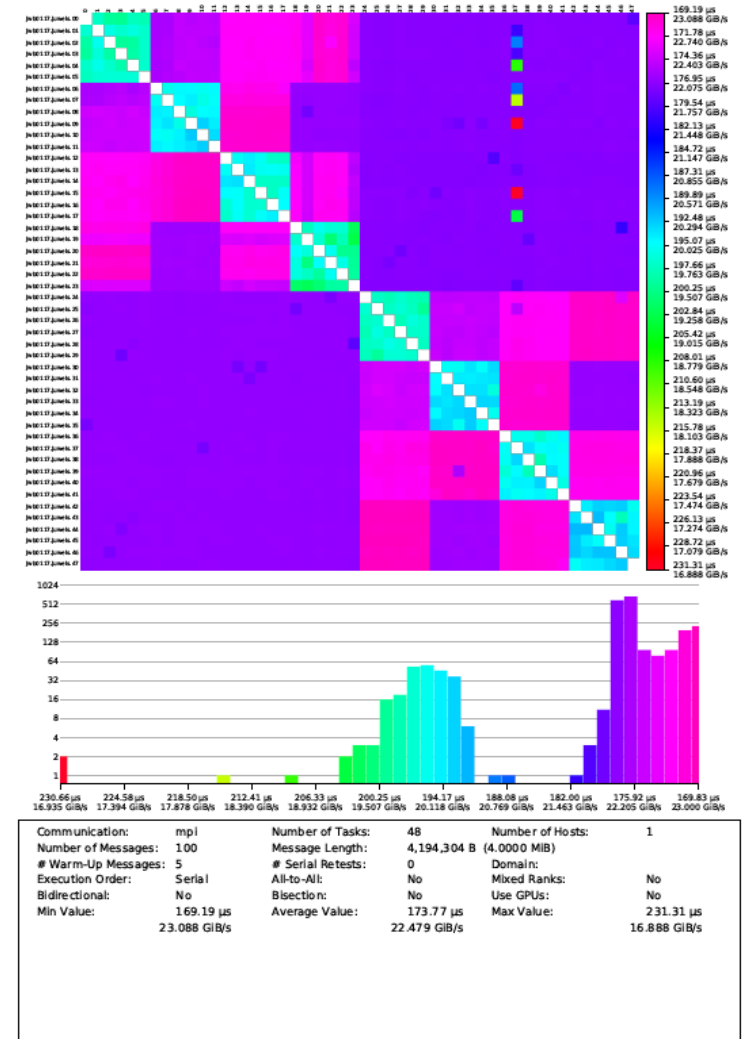


# LINKTEST EXAMPLE

A basic example JUBE benchmark using linktest

# LINKTEST

- Peer-to-peer message-passing timing benchmark
- Supports: MPI, IB verbs, PSM2, UCP, TCP & NVLink via CUDA calls
- Supports (non-)blocking send and recv calls
  - Does not support collective calls aside from MPI all-to-all
- Testing by default in parallel
- Uses either CPU (default) or GPU RAM
- Can perform bidirectional tests
- Can perform bisection tests





# WHAT DOES A SAMPLE JUBE SCRIPT LOOK LIKE?

This is XML!  
YAML also  
works.

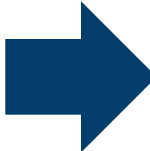
```
1  <?xml version="1.0" encoding="UTF-8"?>
2  <jube>
3    <benchmark name="JSC Linktest Test Suite" outputPath="runs">
4      <comment>Sample Linktest Benchmark</comment>
5
6      <!-- Parameter Sets -->
7      <parameterset name="Linktest Parameters">
9
10     <!-- File Sets -->
11     <fileset name="Linktest Benchmark Files">
13
14     <!-- Substitution Sets -->
15     <substituteset name="Linktest Substitutions">
17
18     <!-- Steps -->
19     <step name="Compile" >
50     <step name="Execute" depend="Compile">
56     <step name="Post Process" depend="Execute">
59
60     <!-- Regex patterns -->
61     <patternset name="Linktest Patterns">
69
70     <!-- Analyse -->
71     <analyser name="Analyser">
76
77     <!-- Results -->
78     <result>
95
96     </benchmark>
97 </jube>
98
```

# WHAT JUBE SETS LOOK LIKE

```
1 <parameterset name="Linktest_Parameters">
2   <parameter name="Number_Of_Nodes" type="int" >4 </parameter>
3   <parameter name="Number_Of_Tasks_Per_Node" type="int" >1 </parameter>
4   <parameter name="Virtual_Cluster_Implementation" type="string" separator="$">mpi$tcp </parameter>
5   <parameter name="Message_Size" mode="python" type="int" >
6     ', '.join(map(str, [2**i for i in range(0,31)]))
7   </parameter>
8   <parameter name="System Name" mode="shell" type="string" >echo -n $$${SYSTEMNAME}</parameter>
9   <parameter name="Date_And_Time" mode="shell" type="string" >date </parameter>
10  <parameter name="Output_Filename_Base" type="string" >
11    linktest_${Virtual_Cluster_Implementation}_${Number_Of_Nodes}nx${Number_Of_Tasks_Per_Node}
12  </parameter>
13 </parameterset>
14
15 <fileset name="Linktest_Benchmark_Files">
16   <copy>Execute_Base.sbatch</copy>
17 </fileset>
18
19 <substituteset name="Linktest_Substitutions">
20   <iofile in="Execute_Base.sbatch" out="Execute.sbatch" />
21   <sub source="$NUMBER_OF_NODES" dest="$${Number_Of_Nodes}" />
22   <sub source="$NUMBER_OF_TASKS_PER_NODE$" dest="$${Number_Of_Tasks_Per_Node}" />
23   <sub source="$VIRTUAL-CLUSTER_IMPLEMENTATION$" dest="$${Virtual_Cluster_Implementation}" />
24   <sub source="$MESSAGE_SIZE$" dest="$${Message_Size}" />
25   <sub source="$OUTPUT-FILENAME_BASE$" dest="$${Output_Filename_Base}" />
26 </substituteset>
```

# HOW SUBSTITUTIONS WORK

```
1 <substituteset name="Linktest_Substitutions">
2   <iofile in="Execute_Base.sbatch" out="Execute.sbatch" />
3   <sub source="$NUMBER_OF_NODESS" dest="{Number_Of_Nodes}" /> <!-- 4 -->
4   <sub source="$NUMBER_OF_TASKS_PER_NODESS" dest="{Number_Of_Tasks_Per_Node}" /> <!-- 1 -->
5   <sub source="$VIRTUAL-CLUSTER_IMPLEMENTATIONS" dest="{Virtual_Cluster_Implementation}" /> <!-- mpi -->
6   <sub source="$MESSAGE_SIZES" dest="{Message_Size}" /> <!-- 1 GiB -->
7   <sub source="$OUTPUT-FILENAME_BASESS" dest="{Output_Filename_Base}" /> <!-- Linktest_mpi_4nxc -->
8 </substituteset>
9 #!/bin/bash
10 #SBATCH --account=cstao
11 #SBATCH --partition=batch
12 #SBATCH --nodes=$NUMBER_OF_NODESS
13 #SBATCH --ntasks-per-node=$NUMBER_OF_TASKS_PER_NODESS
14 #SBATCH --cpus-per-task=1
15
16 # 1. Set Up Arguments
17 args="--mode SVIRTUAL-CLUSTER_IMPLEMENTATIONS \
18 --num-warmup-messages 5 --num-messages 50 \
19 --size-messages $MESSAGE_SIZES \
20 --output $OUTPUT-FILENAME_BASESS.ion"
21
22 # 2. Execute Benchmark
23 srun Compile/linktest/benchmark/linktest ${args};
24
25 # 3. Indicate Success
26 if [ $? -ne 0 ]; then
27   touch error;
28 else
29   touch ready;
30 fi
31
32 exit 0;
```



```
1 #!/bin/bash
2 #SBATCH --account=cstao
3 #SBATCH --partition=batch
4 #SBATCH --nodes=4
5 #SBATCH --ntasks-per-node=1
6 #SBATCH --cpus-per-task=1
7
8 # 1. Set Up Arguments
9 args="--mode mpi \
10 --num-warmup-messages 5 --num-messages 50 \
11 --size-messages 1073741824 \
12 --output Linktest_mpi_4nxc.ion"
13
14 # 2. Execute Benchmark
15 srun Compile/linktest/benchmark/linktest ${args};
16
17 # 3. Indicate Success
18 if [ $? -ne 0 ]; then
19   touch error;
20 else
21   touch ready;
22 fi
23
24 exit 0;
```

# WHAT THE COMPILE STEP MIGHT LOOK LIKE

```
1 <step name="Compile">
2   <do done_file="ready" error_file="error">
3     set -x;
4     # Load Modules
5     module --quiet load Stages/2020;
6     module --quiet load GCC/9.3.0;
7     module --quiet load ParaStationMPI/5.4.7-1;
8     module --quiet load SIONlib/1.7.6;
9     # Clone MiniPMI and Linktest Repositories
10    git clone https://github.com/kraused/minipmi.git;
11    git clone https://deep_bench:_____@gitlab.version.fz-juelich.de/cstao-public/linktest.git;
12    # Compile MiniPMI
13    cd minipmi;
14    make clean;
15    make --jobs=4;
16    LIBRARY_PATH=$(pwd):$(LIBRARY_PATH);
17    CPATH=$(pwd):$(CPATH);
18    # Compile Linktest
19    cd ../linktest/benchmark;
20    make clean;
21    make HAVE_MINIPMI=1 HAVE_IBVERS=0 HAVE_PSM2=0 HAVE_CUDA=0 HAVE_UCP=0;
22    # Check If Compile Succeeded
23    if [[ $? -eq 0 &&& -f "linktest" &&& -f "linktest.mpi" &&& -f "linktest.ibverbs" &&& -f "linktest.tcp" ]];
24    then
25      echo "Linktest compile succeeded" > ../../ready;
26    else
27      echo "Linktest compile failed" > ../../error;
28    fi
29    set +x;
30  </do>
31 </step>
```

Access Token

# WHAT THE OTHER STEPS LOOK LIKE

```
1 <step name="Execute" depend="Compile" >
2   <use >Linktest_Parameters </use>
3   <use >Linktest_Benchmark_Files</use>
4   <use >Linktest_Substitutions </use>
5   <do done_file="ready" error_file="error">sbatch Execute.sbatch </do >
6 </step>
7 <step name="Post_Process" depend="Execute">
8   <do> </do>
9 </step>
```

You can submit batch jobs, just remember to create a file to indicate completion/error. JUBE tests for this to check if it should continue running the benchmark.

The post-process step in this case is not required.

```
1 #!/bin/bash
2 #SBATCH --account=cstao
3 #SBATCH --partition=batch
4 #SBATCH --nodes=4
5 #SBATCH --ntasks-per-node=1
6 #SBATCH --cpus-per-task=1
7
8 # 1. Set Up Arguments
9 args="--mode mpi \
10 --num-warmup-messages 5 --num-messages 50 \
11 --size-messages 1073741824 \
12 --output Linktest_mpi_4nxc.sion"
13
14 # 2. Execute Benchmark
15 srun Compile/linktest/benchmark/linktest ${args};
16
17 # 3. Indicate Success
18 if [ $? -ne 0 ]; then
19   touch error;
20 else
21   touch ready;
22 fi
23
24 exit 0;
```

# HOW RESULTS WORK

<https://regex101.com/>

```
1 <patternset name="Linktest_Patterns">
2   <pattern name="min_time">RESULT: Min Time:\s(${jube_pat_nfp})?\s[n|u|m| |k|M|G|T|P|E|Z|Y]s </pattern>
3   <pattern name="avg_time">RESULT: Avg Time:\s(${jube_pat_nfp})?\s[n|u|m| |k|M|G|T|P|E|Z|Y]s </pattern>
4   <pattern name="max_time">RESULT: Max Time:\s(${jube_pat_nfp})?\s[n|u|m| |k|M|G|T|P|E|Z|Y]s </pattern>
5   <pattern name="min_bw" >RESULT: Min Time:\s(${jube_pat_nfp})?\s[n|u|m| |k|M|G|T|P|E|Z|Y]s \(\s*(${jube_pat_nfp})\s[ |k|M|G|T|P|E|Z|Y][i] ]B/s\) </pattern>
6   <pattern name="avg_bw" >RESULT: Avg Time:\s(${jube_pat_nfp})?\s[n|u|m| |k|M|G|T|P|E|Z|Y]s \(\s*(${jube_pat_nfp})\s[ |k|M|G|T|P|E|Z|Y][i] ]B/s\) </pattern>
7   <pattern name="max_bw" >RESULT: Max Time:\s(${jube_pat_nfp})?\s[n|u|m| |k|M|G|T|P|E|Z|Y]s \(\s*(${jube_pat_nfp})\s[ |k|M|G|T|P|E|Z|Y][i] ]B/s\) </pattern>
8 </patternset>
9
10 <analyser name="Analyser">
11   <analyse step="Execute">
12     <file use="Linktest_Patterns">linktest.log</file>
13   </analyse>
14 </analyser>
15
16 <result>
17   <use>Analyser</use>
18   <table name="result" style="pretty" sort="number">
19     <column title="Benchmark #" >jube_benchmark_id </column>
20     <column title="Date and Time">Date_And_Time </column>
21     <column title="System" >System_Name </column>
22     <column title="Mode" >Virtual_Cluster_Implementation</column>
23     <column title="Message Size" >Message_Size </column>
24     <column title="Min. Time" >min_time </column>
25     <column title="Avg. Time" >avg_time </column>
26     <column title="Max. Time" >max_time </column>
27     <column title="Min. BW" >min_bw </column>
28     <column title="Avg. BW" >avg_bw </column>
29     <column title="Max. BW" >max_bw </column>
30   </table>
31 </result>
```

# HOW RESULTS WORK

## Meta Data

1	Benchmark #	Date and Time	System
2			
3	0	Thu May 20 18:02:17 CEST 2021	jusuf
4	0	Thu May 20 18:02:17 CEST 2021	jusuf
5	1	Thu May 20 18:12:24 CEST 2021	jusuf
6	1	Thu May 20 18:12:25 CEST 2021	jusuf
7	2	Thu May 20 18:14:14 CEST 2021	jusuf
8	2	Thu May 20 18:14:14 CEST 2021	jusuf

## Benchmark Data


Mode	Message Size	Min. Time
mpi	1	1.66078098 us
tcp	1	50.63614808 us
mpi	1	1.65038276 us
tcp	1	49.89886191 us
mpi	1	1.63517892 us
tcp	1	53.37819923 us

1	Avg. Time	Max. Time	Min. BW	Avg. BW	Max. BW
2					
3	1.70571070 us	1.80307776 us	588.014 kiB/s	572.525 kiB/s	541.609 kiB/s
4	59.48583324 us	73.48540239 us	19.286 kiB/s	16.417 kiB/s	13.289 kiB/s
5	1.71088614 us	1.93268061 us	591.719 kiB/s	570.793 kiB/s	505.289 kiB/s
6	69.69922998 us	97.96963073 us	19.571 kiB/s	14.011 kiB/s	9.968 kiB/s
7	1.66639569 us	1.70407817 us	597.221 kiB/s	586.033 kiB/s	573.074 kiB/s
8	75.47586438 us	111.84767820 us	18.295 kiB/s	12.939 kiB/s	8.731 kiB/s

# THE JUBE CRON JOB

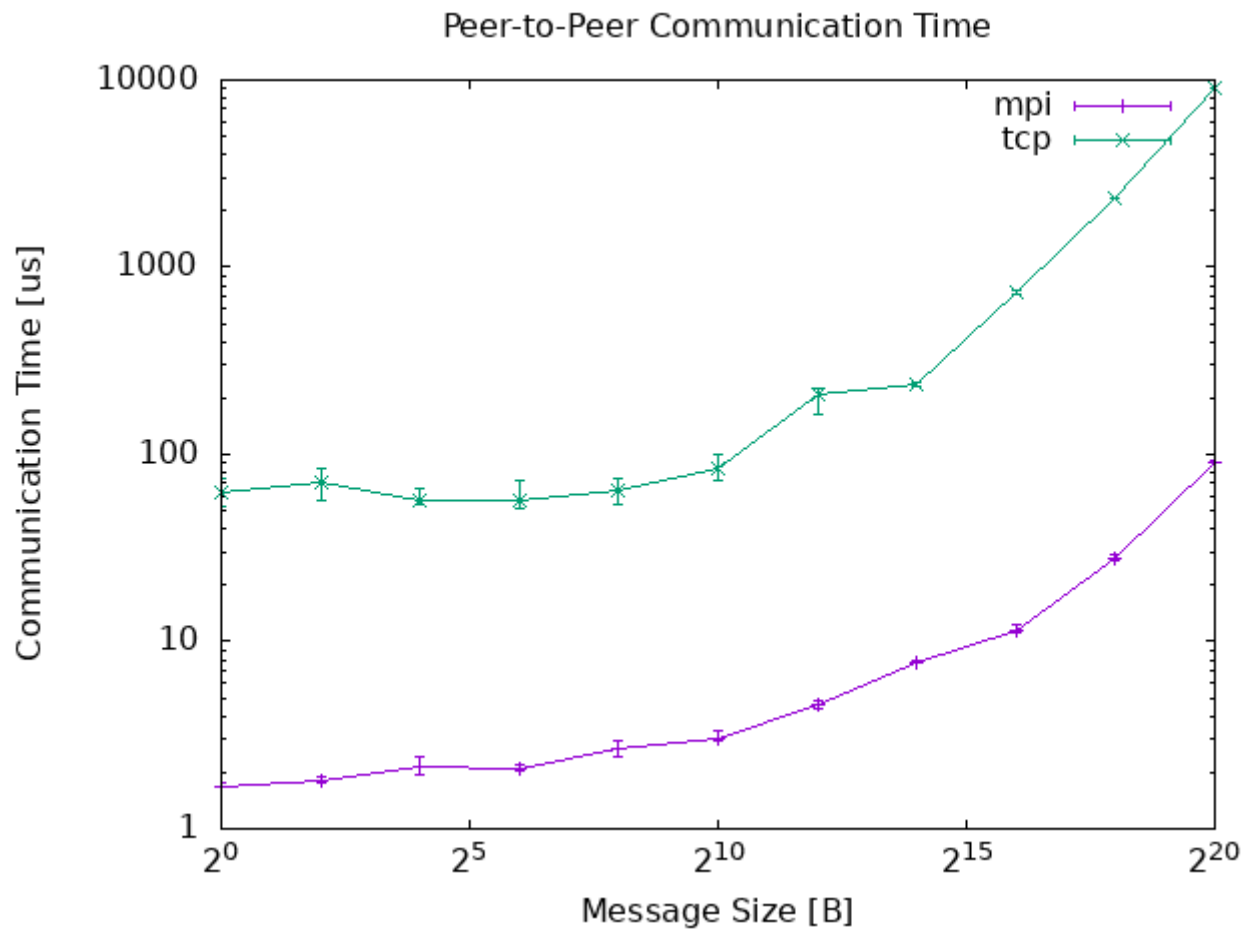
```
1  #!/bin/bash
2
3  # 1. Load Modules
4  module --quiet load JUBE/2.4.1;
5
6  # 2. Execute Benchmark
7  jube-autorun -o Linktest.xml >/dev/null;
8
9  # 3. Get Result Table
10 mkdir -p Results;
11 Date_Of_Year="$(date +%d-%m-%Y)";
12 Filebase="Results/Linktest_Size_Test";
13 Filename="${Filebase}_${Date_Of_Year}.log";
14 jube result runs | tail -n +2 | head -n -1 >${Filename};
15
16 # 4. Append Results To Master Table
17 Mastername="${Filebase}.log";
18 if [ ! -f ${Mastername} ]; then cp ${Filename} ${Mastername};
19 else tail -n +3 <${Filename} >>${Mastername}; fi;
20
21 # 5. Upload Result Table
22 git checkout ${SYSTEMNAME};
23 git add ${Filename} ${Mastername};
24 git commit -m "Results for ${Date_Of_Year}";
25 git push;
26
27 # 6. Archive Benchmark
28 sleep 30;
29 for dir in $(find runs ! -path '*/result' -mindepth 2 -maxdepth 2 -type d)*/;do
30     tar -cjf ${dir}.tar.bz2 ${dir} && rm -rf ${dir};
31 done
32
33 exit;
```

Insert result-table  
post-processing here.





# POST-PROCESSING EXAMPLE



# ONE LAST TIP

You can use parameters to select other parameters!

```
1 <parameter name="a" mode="shell" type="int" >1 </parameter>
2 <parameter name="b" mode="shell" type="int" >2 </parameter>
3 <parameter name="c" mode="shell" type="int" >3 </parameter>
4 <parameter name="i" mode="shell" type="int" >0,1,2 </parameter>
5 <parameter name="X" mode="python" type="int" >[ 1 , 2 , 3 ][${i}]</parameter>
6 <parameter name="Y" mode="python" type="string">["a" ,"b" ,"c" ][${i}]</parameter>
7 <parameter name="Z" mode="python" type="int" >[${a},${b},${c}][${i}]</parameter>
```



# Thank You! Questions?!

DR. MAX HOLICKI | JÜLICH SUPERCOMPUTING CENTRE



# Live Demonstration

DR. MAX HOLICKI | JÜLICH SUPERCOMPUTING CENTRE

# EFFECTIVE JUBE BENCHMARK CREATION

- KISS (Keep It Simple Stupid)
- Use precompiled binaries
- Skip compile step for now (replace it with a copy)
- Only include relevant parameters at the start
- Use Python for parameters
- Use JUBE debug mode for dry runs 1<sup>st</sup>
- To set up the regex patterns use update (saves rerunning)
- Use a regex pattern builder
- Do not forget the ready/error files
- Have a look at the JUBE glossary

<https://apps.fz-juelich.de/jsc/jube/jube2/docu/glossar.html>

# Have Fun!