

Fair-share policy on IAP compute servers

Dennis Arslan
Daniel Repp
Katsuya Tanaka
Maximilian Weißflog

16th August, 2022

Contents

1	General rules of conduct	1
1.1	Logging out versus disconnecting from a remote desktop session	1
1.2	Keeping the local drives clean	2
1.3	Parallel computing	4
2	COMSOL	5
2.1	Folder settings	5
2.2	Limiting the number of processors	5
3	Lumerical	7
3.1	Limiting the number of processors	7
4	MATLAB	10
4.1	Limiting the number of processors	10
5	Thales – Advanced topics	11
	List of Figures	12

1 General rules of conduct

Our IAP compute servers, such as Heron and Thales, are used by a large number of group members for computation-intensive tasks. However, since there is no automatic scheduling of resources, it is the responsibility of each user to ensure that their processes are not overusing resources or negatively interfering with the processes of other users. Misconduct can not only slow down the entire server, but may even lead to severe crashes. In this document, we want to make you aware of the most important issues and advise you on how to deal with them.

For each type of resource, there are a few critical aspects to consider:

CPU If there are more concurrent long-running threads than there are physical processors, the entire server may freeze. Each physical processor can work on only one thread at a time. Make sure that your programs spawn only a limited number of threads and leave enough processors for other users. On Thales, we also bind processes to specific processors, so that users can use them exclusively. This usually results in a small performance increase and processes of different users interfere less with each other.


RAM Currently, it seems that we hardly ever come close to the capacity limit. For good practice, however, you should always release the RAM once you are done with your computations.

Local drives The operating system and most programs are installed on the C-drive. If this drive runs out of memory, various programs of every user on the server may crash and certain files can get corrupted. Simply do not store any private data on the C-drive, except when it is inevitable for the operation of your program. Unfortunately, most programs put their temporary folders on the C-drive by default. If your program produces large amounts of temporary data, you must move its temporary folder to another local drive. Every user must regularly remove their private data from all local drives and delete any residual temporary files which were not automatically deleted by the respective program.


Network bandwidth If your program constantly reads or writes large amounts of data, make sure that these files are on a local drive and not the network drive.

1.1 Logging out versus disconnecting from a remote desktop session

Problem Disconnecting from a remote desktop session does not log you out, so that you can continue your work from where you left. However, even if your programs are in an idle state, every logged-in user consumes a small amount of RAM and CPU-time. This quickly sums up with our large number of users. On Heron, we even observed up to 30% CPU utilization when all users were idle. Figure 1.1a, for example, shows all the disconnected users (Status: Getrennt) and their occupied resources.



(a) Task Manager, showing the state and resources of each user.



(b) Start screen, showing how to log out from the remote desktop session.

Figure 1.1: How to log out from a remote desktop session on Heron and Thales.

Solution

1. **Whenever you are done with your work on the server, completely log out from your user account in order to free up resources.**

If you think about the server as if it was your own desktop PC, then logging out is effectively the same as shutting down and disconnecting is the same as turning off the screen.

Figure 1.1b shows how to log out from Heron and Thales. Go to “Windows Start”, click on your user name, then click on “Abmelden”.

1.2 Keeping the local drives clean

Problem When the C-drive runs out of memory and any program of any user tries to write to that drive, it will most likely crash. This usually results in aborted simulations, loss of data, or even corrupted system files. We encountered cases where Matlab would not start up again because of corrupted configuration files. The entire server can become unusable until the person who caused the error takes some action. Therefore, it is critical to keep enough free space on the C-drive. When the other local drives run out of memory, we may still encounter aborted simulations and loss of data, but at least no corrupted system files. Consequently, it is important to keep enough free space on all local drives.

Solution

1. **Do not store any big files on the C-drive.**

Let’s estimate the maximum space on Heron’s C-drive you are theoretically allowed to occupy: There are ca. 70 user profiles, the C-drive has a total capacity of 446 GB, but the operating system and other programs already occupy around 136 GB. This leaves us with around 4.5 GB for every user. However, this is really the best-case scenario which would still lead to a full drive. In this situation, it may be acceptable if you occupied up to about 2 or 3 GB. However, we also believe that there is usually no reason for any user to occupy more than 1 GB. Note that these limits include both private and temporary data.

2. **Ideally, move all your temporary folders to another local drive.**

On Heron and Thales, the other local drives are SSDs with more capacity than the C-drive. If you run programs which generate huge amounts of temporary data, you must make sure that the temporary folders are on the other local drives. Since all our drives are currently pretty full, you may keep certain temporary folders on the C-drive if you can guarantee that they will not take up much space.

3. **Regularly delete your private data or move it off the server to a permanent storage location.**

4. **Regularly delete any temporary files which were not automatically deleted by your programs.**

Depending on the programs you use, there can be several locations where your temporary files are hiding. In the following, we give some hints on where to look and what to do.

Start by checking your Windows user folder (%USERPROFILE%) for large files and folders, then dig down the folder tree. Also make sure that you look inside the app data folder (%USERPROFILE%\AppData), which may be hidden by default. It is likely that all your data on the C-drive is contained within these folders. You can quickly access these folders via the Windows Explorer (see fig. 1.2).

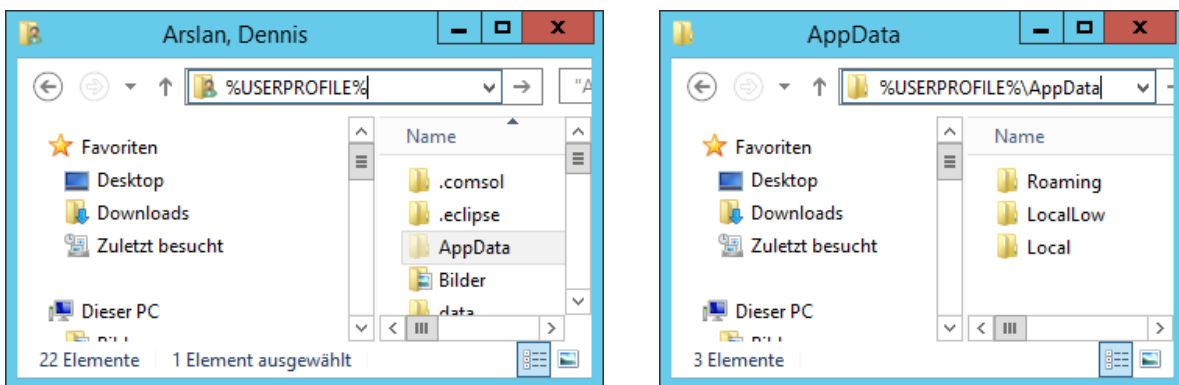


Figure 1.2: Folders on the C-drive which usually contain all user-specific data.

You can create a new temporary folder on another drive and change the environment variables **TEMP** and **TMP** such that they point to this new folder. Note that this is the default temporary folder which Windows offers to all programs. Be aware that some programs may not use this folder, or even create their own temporary folder somewhere else. Figure 1.3 shows how I moved the Windows temporary folder to D:\users\arslan\tmp.

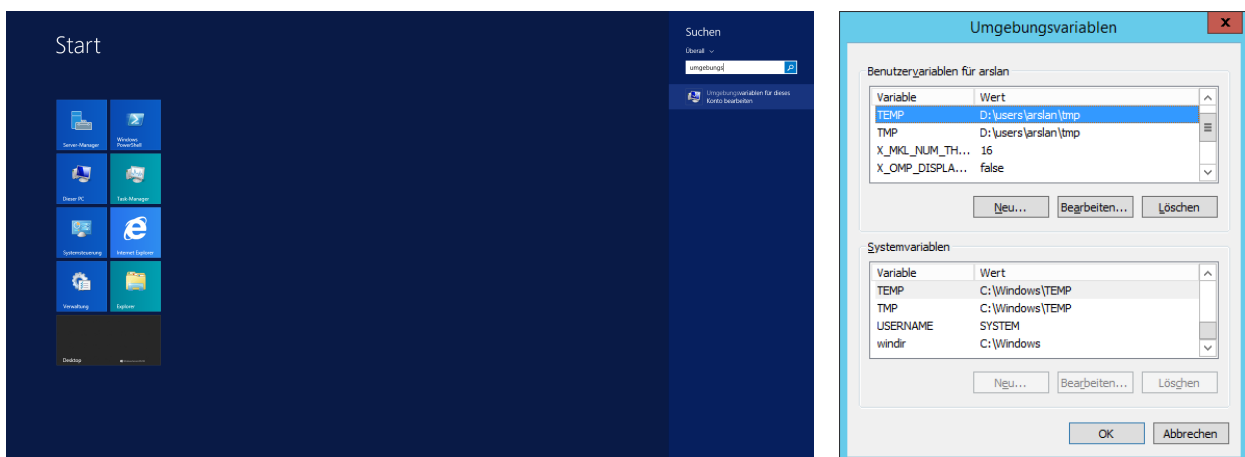


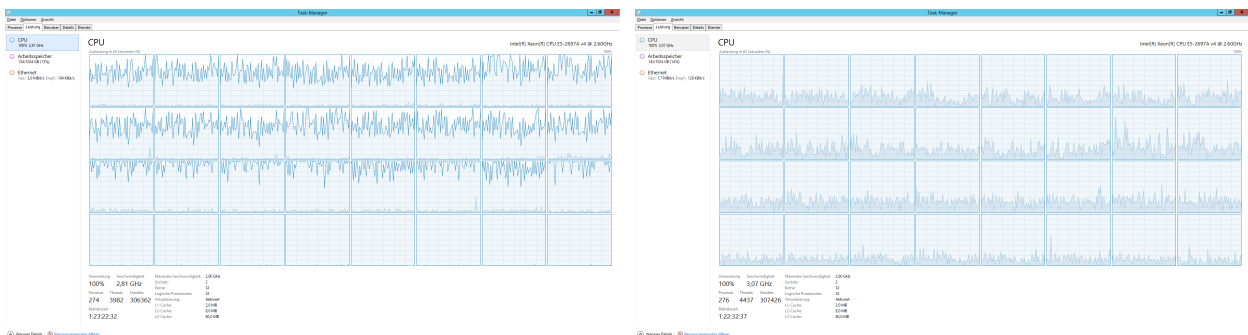
Figure 1.3: How to move the Windows temporary folder location by editing the TEMP and TMP environment variables.

1.3 Parallel computing

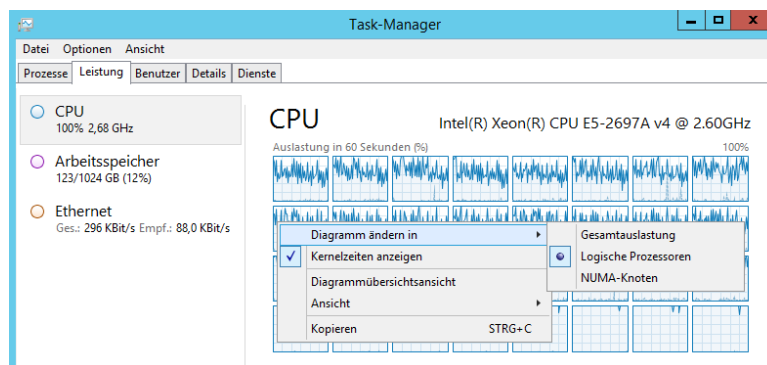
Problem Depending on how well your program can parallelize its work, there is usually an optimal number of processors for which your program will run the fastest (Amdahl's law¹). In most cases, you will even observe a slow-down if you use too many processors. This means that these additional processors are not only useless or even detrimental for your program, but they are also not available for other users.

Solution Usually, you will have to determine the optimal number of processors via trial and error. Sometimes you can run a small part of your workload and compare the execution time for different numbers of processors. At other times, it can be sufficient to look at the Task Manager.

Figure 1.4c shows how you can configure the Windows Task manager to show the utilization of each processor and the kernel time. The kernel time tells you how much time the processor spends on low-level tasks of the operating system, as for example communicating with or waiting for other processors. If the kernel time is constantly high, it usually means that there is something wrong with your settings.



(a) Almost full utilization of all processors with almost no kernel time (dark blue). (b) Full utilization of all processors, but roughly 20% of the time is wasted in the kernel (dark blue).



(c) How to configure the display.

Figure 1.4: How to interpret the CPU performance tab in the Windows Task Manager.

¹We suggest that you make yourself familiar with the basic concepts of parallel computing to better understand the implications of our suggestions. For example, see: https://en.wikipedia.org/wiki/Parallel_computing#Background

2 COMSOL

2.1 Folder settings

Figure 2.1 summarizes some optimal settings regarding the location of temporary files.

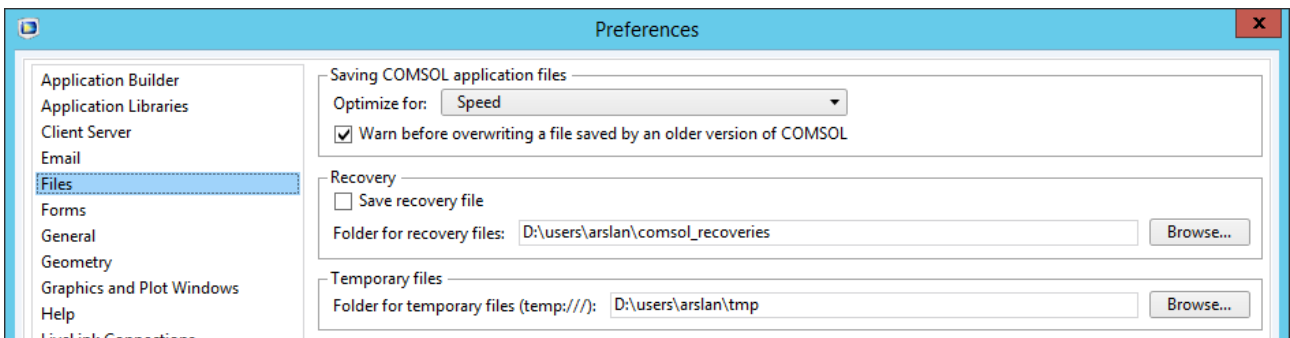


Figure 2.1: COMSOL preferences “Files”

Save recovery file Disable recovery files to make your simulations complete faster and produce less junk data! These recovery files are hardly ever useful. Especially if your model is big, COMSOL will waste a lot of extra time and memory on writing these recovery files.

Folder for recovery files Create a new folder on a local drive (not on C:) and reference it here. Due to some reason, COMSOL temporarily writes simulation results into the recovery folder instead of the temporary folder.

Folder for temporary files Create a new folder on a local drive (not on C:) and reference it here. This folder usually contains files which are required by the graphic user interface or COMSOL server.

2.2 Limiting the number of processors

Figure 2.2 shows how to set the default number of processors used by each COMSOL process that you start. Unchecked boxes use the grayed-out default values. Check a box to provide your own settings. Note that these settings usually only apply after a restart of COMSOL.




Figure 2.2: COMSOL preferences for “Multicore and Cluster Computing”

Number of cores This is the number of processors which will be used by each COMSOL process that you start. **WARNING: While this is the normal way to set the number of processors, it will have no or even negative effects on Heron and Thales if you set the number to anything other than 4. See below for an alternative.**

Number of sockets Keep it unchecked unless you know what you are doing.

Memory allocator Keep it unchecked unless you know what you are doing.

The best way to control how many and which processors shall be used, is to start COMSOL with a few command line arguments. We prepared Windows batch scripts for Thales which not only limit the number of processors, but also set some critical environment variables and bind COMSOL to specific processors (see chapter 5). Please do not use the scripts on Heron, since the processor binding can have negative effects there. Furthermore, due to the large number of users on Heron, we also do not see the need to provide modified scripts for Heron.

3 Lumerical

3.1 Limiting the number of processors

If you are on Thales, first communicate with the other users which processors you want to use (see fig. 3.1 and chapter 5). You will have to bind Lumerical to these processors. On Heron, you should not use processor binding.

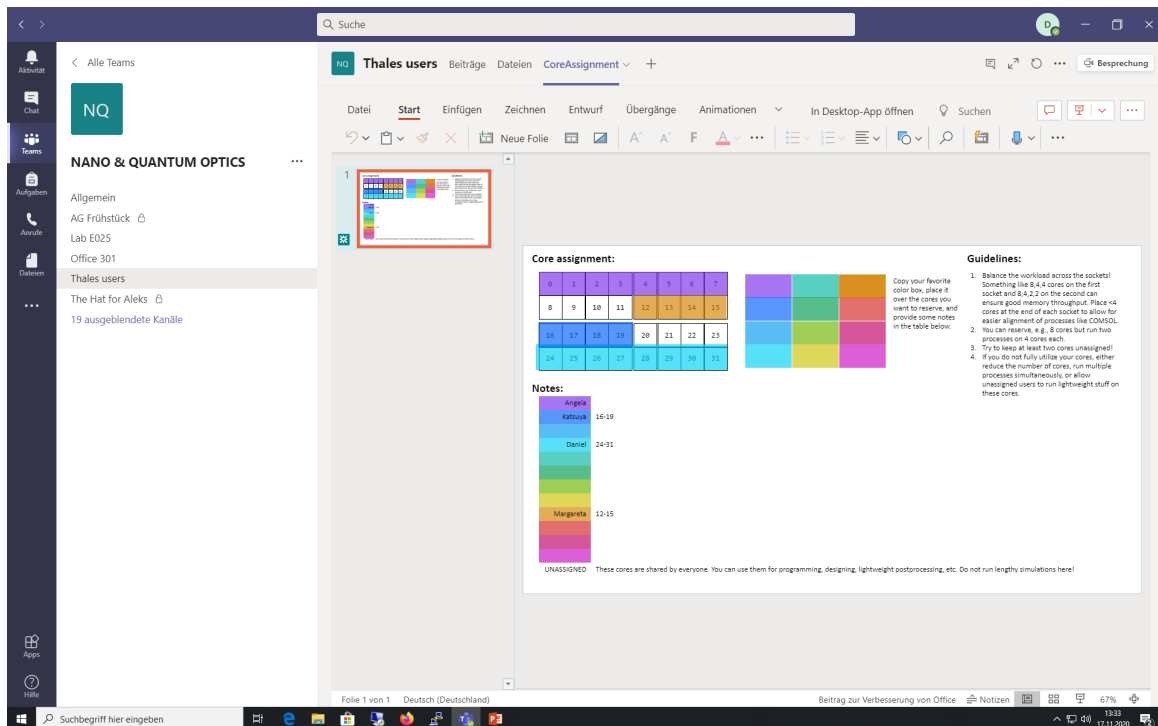


Figure 3.1: MS Teams – Thales users. How to communicate exclusive processor reservation.

Processor binding via the Windows Task Manager For this method to work, you first have to start the graphic user interface of Lumerical, then open the Windows Task manager, go to

Windows Task Manager > Details > `fddt-solutions.exe` > Zugehörigkeit festlegen

and select the cores you reserved (see fig. 3.2). This sets the so-called affinity mask which is managed by Windows. Every child process started by `fddt-solutions.exe` will inherit the affinity mask and run on the same selected processors. The graphic user interface usually starts several `fddt-engine.exe` processes which work in parallel on the simulation. Please verify that these processes are indeed bound to the selected cores.

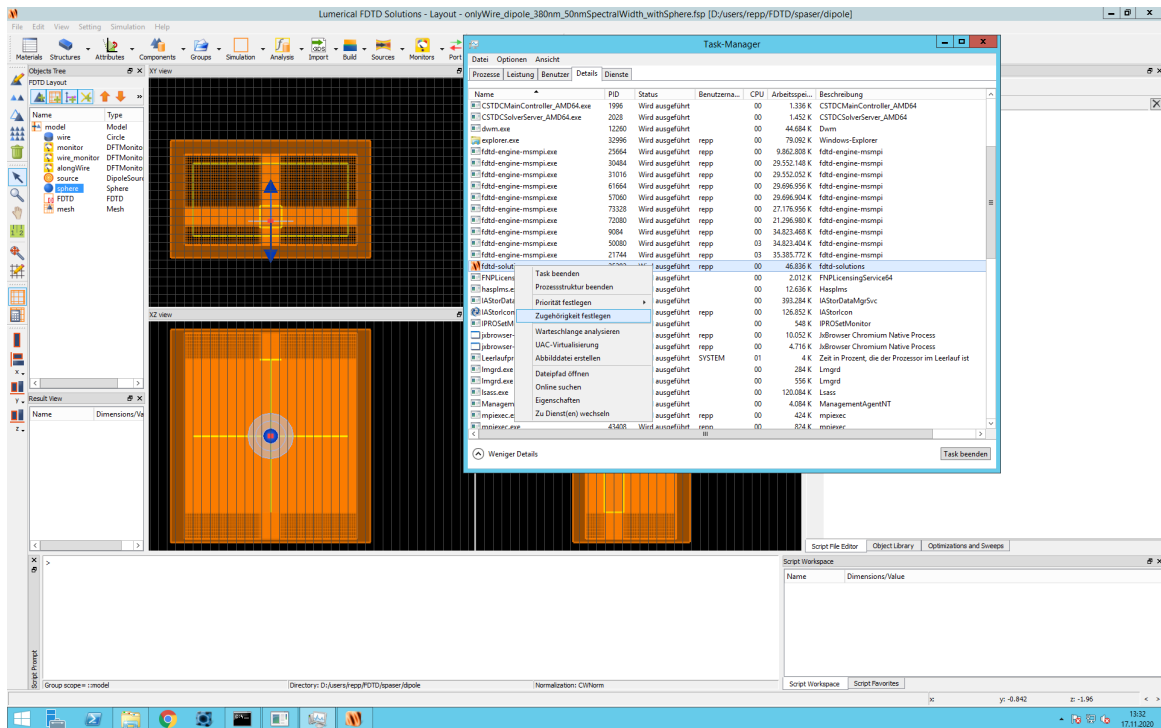


Figure 3.2: How the set the affinity mask for Lumerical.

Processor binding via Windows batch scripts This method allows you to start any Lumerical process with a given affinity mask from the command line or a Windows batch script. The command takes the following form

```
start "" /affinity 0x30000000
"C:\Program Files\Microsoft MPI\Bin\mpiexec.exe" -n 4
"C:\Program Files\Lumerical\FDTD\bin\fddt-engine-msmpi.exe" -t 1
"yourModelFile.fsp"
```

where the line breaks are only for readability and each element has the following meaning:

start <title> /affinity <hex> <program> <parameter> The string <title> will be displayed in the title of the command prompt. Although the documentation says it is an optional argument, you usually have to provide an empty string "" for the command to work correctly. The /affinity <hex> key provides the affinity mask as a hexadecimal number. A detailed description of this number is given below. The sequence <program> <parameter> specifies the program which shall be started with the given affinity mask and the parameters which shall be passed to the program during startup. In the example above, <program> represents the string containing `mpiexec.exe` and <parameter> the remainder.

mpiexec <mpiOptions> <solver> <solverOptions> The `mpiexec.exe` process manages the parallel execution of a particular Lumerical `solver`. In the example above, <mpiOptions> represents `-n 4`, <solver> the string containing `fddt-engine-msmpi.exe` and <solverOptions> the remainder. The option `-n 4` means that up to 4 solver processes are allowed to run in parallel.

fddt-engine <options> <filename> The `fddt-engine-msmpi.exe` performs the simulation. In the example above, <options> represents `-t 1` and <filename> the model file `"yourModelFile.fsp"`. The option `-t 1` implies that each solver process runs on one processor.

Note that you may simply omit the `/affinity <hex>` key to avoid the processor binding. For a full list of all available options, please refer to the respective documentations:

- `start` – Starts a separate Command Prompt window to run a specified program or command.
- `mpiexec` – Starts one or more Message Passing Interface (MPI) applications on an HPC cluster.
- `fttd-engine` – Running simulations via the MPI.

Determining the hexadecimal value of the affinity mask The affinity mask is based on a binary representation of the selected processors. In the following example, we assume that the CPU has 32 processors with identification numbers (IDs) ranging from 0 to 31. We want to select the processors with ID 28 and 29:

Processor ID	31	24	23	16	15	8	7	0	
Binary mask	0011	0000	0000	0000	0000	0000	0000	0000	
	----	----	----	----	----	----	----	----	
Hexadecimal	3	0	0	0	0	0	0	0	= 0x30000000

You may omit the leading zeros. For example, selecting processors 8 to 11 yields:

Processor ID	31	24	23	16	15	8	7	0	
Binary mask	0000	0000	0000	0000	0000	1111	0000	0000	
	----	----	----	----	----	----	----	----	
Hexadecimal	0	0	0	0	0	f	0	0	= 0xf00

Figure 3.3 shows how you can use the Windows calculator tool to quickly find the hexadecimal value of the affinity mask by hand. Simply switch to the programmer mode and click on the bits in the binary representation which you want to set (directly below the output field).

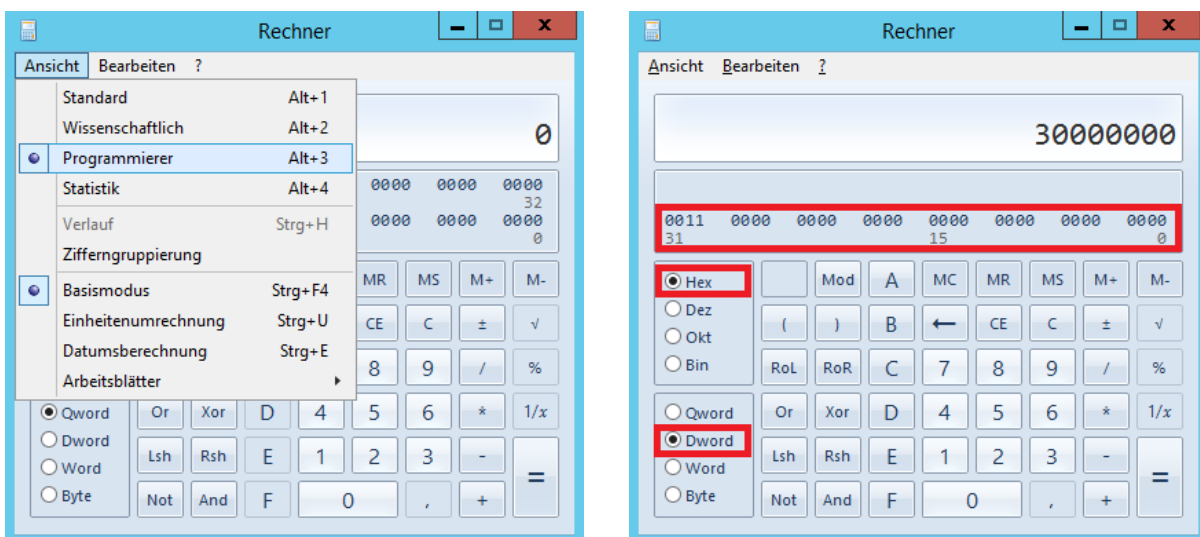


Figure 3.3: Using the Windows calculator to quickly find the affinity mask by hand.

4 MATLAB

4.1 Limiting the number of processors

In MATLAB, you can set the number of computation threads at runtime via the function:

```
maxNumCompThreads(N)
```

See <https://de.mathworks.com/help/matlab/ref/maxnumcompthreads.html> for more information.

It is also possible to bind all MATLAB computation threads to specific processors at runtime by setting the affinity mask via the .NET interface. See chapter 5 for more information.

5 Thales – Advanced topics

On Thales, we agreed on binding all processes to specific processors. This often not only speeds up simulations, but also prevents users from interfering with each other. If you are working on Thales or want learn more about some technical details not covered in this document, join the [Thales users](#) channel on Microsoft Teams. There, you can also find some further information and utility scripts in the “Files” section.

List of Figures

1.1	How to log out from a remote desktop session on Heron and Thales.	2
1.2	Folders on the C-drive which usually contain all user-specific data.	3
1.3	How to move the Windows temporary folder location by editing the TEMP and TMP environment variables.	3
1.4	How to interpret the CPU performance tab in the Windows Task Manager.	4
2.1	COMSOL preferences "Files"	5
2.2	COMSOL preferences for "Multicore and Cluster Computing"	5
3.1	MS Teams – Thales users. How to communicate exclusive processor reservation.	7
3.2	How to set the affinity mask for Lumerical.	8
3.3	Using the Windows calculator to quickly find the affinity mask by hand.	9