



Which Hardware to Buy

When using Avizo®/Amira® to visualize and process large data

Introduction

This document is intended to give recommendations about choosing a suitable workstation to run Avizo/Amira.

The four most important components that need to be considered are the graphics card (GPU), the CPU, the RAM and the hard drive.

The performance of direct volume rendering of large volumetric data or large triangulated surface visualization extracted from the data depends heavily on the GPU capability. The performance of image processing algorithms depends heavily on the performance of the CPU. The ability to quickly load or save large data depends heavily on the hard drive performance. And, of course, the amount of available memory in the system will be the main limitation on the size of the data that can be loaded and processed.

Because the hardware requirements will widely vary according to the size of your data and your workflow, we strongly suggest that you take advantage of our supported evaluation version to try working with one of your typical data sets.

In this document, the term “Avizo” refers to all Avizo editions (including Avizo ToGo) and all Avizo extensions, and the term “Amira” refers to all Amira editions and extensions.

Operating systems

- Microsoft Windows 7/8/10 (64-bit).
- Linux x86 64 (64-bit). Supported 64-bit architecture is Intel64/AMD64 architecture. Supported Linux distribution is Red Hat Enterprise Linux 6.
- Mac OS X 10.9 (64-bit).

Some of the editions and extensions are limited to specific platforms:

- Avizo Inspect is supported only on Microsoft Windows (64-bit).
- Avizo XWind Extension: support of ABAQUS reader (.odb format) and STAR-CCM reader is available only on Microsoft Windows and Linux, not on Mac OS X.
- XScreen is supported only on Microsoft Windows and Linux, not Mac OS X.
- XReadIGES, XReadSTEP, XReadCATIA5 are supported only on Microsoft Windows, not on Linux or Mac OS X.
- XLab Hydro is supported only on Microsoft Windows 64-bit, Linux and Mac OS X, not on Microsoft Windows 32-bit.
- XLab Diffusion, XLab Electro, XLab Thermo are supported only on Microsoft Windows 64-bit, not on Microsoft Windows 32-bit, Linux, or Mac OS X.

Prioritizing hardware for Avizo/Amira

Graphics Card

The single most important determinant of Avizo/Amira performance for visualization is the graphics card.

Avizo/Amira should run on any graphics system (this includes GPU and its driver) that provides a complete implementation of OpenGL 2.1 or higher (certain features may not be available depending on the OpenGL version and extensions supported). However, graphics board and driver bugs are not unusual.

The amount of GPU memory needed depends on the size of the data. We recommend a minimum of 1 GB on the card. Some visualization modules may require having graphics memory large enough to hold the actual data. High-end graphics cards have 12 to 16 GB of memory. Optimal performance volumetric visualization at full resolution requires that data fit in graphics memory (some volume rendering modules of Avizo/Amira are able to go around this limitation).

Avizo/Amira will not benefit from multiple graphics boards for the purpose of visualization on a single monitor. However, some of the image processing algorithms rely on CUDA for computation, and while the computation can run on the single CUDA-enabled graphics board, this computation can also run on a second CUDA-enabled graphics card installed on the system. A multiple graphics board configuration can be useful to drive many screens or in immersive environments.

When comparing graphics boards, there are many different criteria and performance numbers to consider. Some are more important than others, and some are more important for certain kinds of rendering. Thus, it's important to consider your specific visualization requirements. Integrated graphics boards are not recommended for graphics-intensive applications such as Avizo/Amira except for basic visualization.

Wikipedia articles on NVIDIA GeForce/Quadro and AMD Radeon/FirePro cards will detail specific performance metrics:

- Memory size: This is very important for volume visualization (both volume rendering and slices) to maximize image quality and performance because volume data is stored in the GPU's texture memory for rendering. It is also important for geometry rendering if the geometry is very large (large number of triangles).
- Memory interface / Bandwidth: This is important for volume rendering because large amounts of texture data need to be moved from the system to the GPU during rendering. The PCI Express 3 buses are the fastest interfaces available today.
- Number of cores (also known as stream processors): This is very important for volume rendering because every high-quality rendering feature you enable requires additional code to be executed on the GPU during rendering.
- Triangles per second: This is very important for geometry rendering (surfaces, meshes).
- Texels per second / Fill rate: This is very important for volume visualization (especially for volume rendering), because a large number of textures will be rendered and pixels will be "filled" multiple times to blend the final image.

Professional graphics boards

Vendor	Family	Series
NVIDIA	Quadro	Maxwell, Kepler, Fermi
AMD	FirePro	W, V

All driver bugs are submitted to the vendors. A fix may be expected in a future driver release.

Standard graphics boards

Vendor	Family	Series
NVIDIA	GeForce	Maxwell, Kepler, Fermi
AMD	Radeon	since Evergreen
Intel	HD Graphics	Ivy Bridge, Haswell

Due to vendor support policies, on standard graphics boards we are not able to commit to providing a fix for bugs caused by the driver.

1. A professional graphics boards will benefit from the professional support offered by the vendors (driver bug fixes)
2. Always use a recent driver version for your graphics board.
3. With an NVIDIA Quadro board we recommend to use the driver profile "3D App - Visual Simulation". In case of rendering or performance issues you may want to experiment with different "3D App" profiles.
4. Visit http://en.wikipedia.org/wiki/List_of_Nvidia_graphics_processing_units for a complete list of NVIDIA boards and comparisons.
5. Visit http://en.wikipedia.org/wiki/List_of_AMD_graphics_processing_units for a complete list of AMD boards and comparisons.

System Memory

System memory is the second most important determinant for Avizo/Amira users who need to process large data.

You may need much more memory than the actual size of the data you want to load within Avizo/Amira. Some processing may require several times the memory required by the original data set. If you want to load, for instance, a 4 GB data set in memory and apply a non-local means filter to the original data and then compute a distance map, you may need up to 16 or 20 GB of additional memory for the intermediate results of your processing. Commonly you will need 2 or 3 times the memory footprint of the data being processed for basic operations. For more complex workflows you may need up to 6 or 8 times amount of memory, so 32 GB may be required for a 4 GB dataset. Due to the potentially high memory requirement we highly recommend running on 64-bit computer and operating systems.

Also notice that size of the data on disk may be much smaller than memory needed to load the data as the file format may have compressed the data (for instance, loading a stack of JPEG files).

Avizo's/Amira's Large Data Access (LDA) technology will enable you to work with data sizes exceeding your system's physical memory. LDA is an excellent way to stretch the performance, but it is not a direct substitute for having more physical memory.

The best performance and optimal resolution will be achieved by using Avizo's/Amira's LDA technology in combination with a large amount of system memory. LDA provides a very convenient way to quickly load and browse your whole dataset. Note that LDA data will not work with most compute modules, which require the full resolution data to be loaded in memory.

Avizo and Amira XImagePAQ provide another loading option to support 2D and 3D image processing from disk to disk, without requiring loading the entire data into memory; modules then operate per data slab. This enables processing and quantification of large image data even with limited hardware memory. Since processing of each slab requires loading data and saving results from/to the hard drive, it dramatically increases processing time. Thus, processing data fully loaded in memory is always preferred for best performance.

Hard Drives

When working with large files, reading data from the disk can slow down your productivity. A standard hard drive (HDD) (e.g., 7200rpm SATA disk) can only stream data to your application at a sustained rate of about 60 MB/second. That is the theoretical limit; your actual experience is likely to be closer to 40 MB/second. When you want to read a 1 GB file from the disk, you will likely have to wait 25 seconds. For a 10 GB file, the wait is 250 seconds, over 4 minutes. LDA technology will greatly reduce wait time for data visualization, but disk access will still be a limiting factor when you want to read data files at full resolution for data processing. Compared to traditional HDDs, solid state drives (SSD) can improve read and write speeds.

For best performance, the recommended solution is to configure multiple hard drives (3 or more HDD or SSD) in RAID5 mode; note that RAID configurations may require substantially more system administration. For performance only, RAID 0 could be used, but be warned of risk of data loss upon hard-drive failure. If you want performance and data redundancy then RAID 5 is recommended.

Reading data across the network, for example from a file server, will normally be much slower than reading from a local disk. The performance of your network depends on the network technology (100 Mb, 1 Gb, etc.), the amount of other traffic on the network, and number/size of other requests to the file server. Remember, you are (usually) sharing the network and server and will not get the theoretical bandwidth. LDA technology may also facilitate visualization of volume data through the network, but if data loading is a bottleneck for your workflow, we recommend making a local copy of your data.

CPU

While Amira/Avizo mostly relies on GPU performance for visualization, many modules are computational intensive and their performance will be strongly affected by CPU performance.

More and more modules inside Avizo/Amira are multi-threaded and thus can take advantage of multiple CPUs or multiple CPU cores available on your system. This is the case for most of the quantification modules provided with Avizo and Amira XImagePAQ, a number of modules of the XLab Extension, and also various computation modules.

Fast CPU clock, number of cores, and memory cache are the three most important factors affecting Avizo and Amira performance. While most multi-threaded modules will scale up nicely according to the number of cores, the scaling bottleneck may come from memory access. From experience, up to 8 cores show almost linear scalability while more than 8 cores do not show much gain in performance. A larger memory cache improves performance.

Special considerations regarding Avizo ToGo

All hardware recommendations listed above are also relevant for Avizo ToGo when loading a .togo file generated by Avizo.

The amount of memory needed will be greater than the total memory footprint of the data stored in the .togo file plus the memory footprint of the application itself. Note, since data are compressed when stored in the .togo file and the compression rate varies hugely depending on the data, there is no way to predict from the file size the amount of memory needed to load a .togo file.

Graphics recommendations are the same as above, and will vary depending on the complexity of the geometry to be displayed and the size of the data to be rendered with volume rendering.

The CPU recommendations are less demanding than above, as .togo only stores the result of computation and does not run the computation itself.

It is highly recommended before exporting a .togo file from Avizo to remove all unnecessary data and modules in order to minimize the size of the file and also the amount of memory that will be needed by Avizo ToGo. In order to get an estimate of the required system memory for the workstation where the Avizo ToGo file will be loaded, you can use the Task Manager or other system utilities to look at the current Avizo memory usage at the time of export. You may want to provide this information to the recipient of the .togo file so he or she can provide appropriate hardware resources.

How hardware can help optimizing

Here is a summary of hardware characteristics to consider for optimizing particular tasks.

Visualizing large data (LDA)

- Fast hard drive
- System memory
- GPU memory
- Memory to GPU/CPU bandwidth

Basic volume rendering

- GPU fill rate (texels per second)

Advanced volume rendering (Volume Rendering module):

- Heavy use of pixel shaders
- GPU clock frequency, number of GPU cores

Large geometry rendering such as large surfaces from Isosurface or Generate Surface, large point clusters, large numerical simulation meshes...

- GPU clock frequency, number of triangles per second

Image processing and quantification (Avizo, Amira XImagePAQ)

- Multiple CPU cores (for many modules, including most image processing modules)
- CPU clock frequency

Anisotropic Diffusion, Non-Local Means Filter (high-performance smoothing and noise reduction image filters)

XLab Hydro (absolute permeability computation)

- GPU speed, number of GPU cores (stream processors), CUDA-compatible (NVIDIA)

Other compute modules, display module data extraction

- CPU clock frequency
- Multiple CPU cores (for a number of multi-threaded modules, such as Generate Surface, Register Images, Resample, Arithmetic)

Multi-display systems with Avizo/Amira XScreen: tiled displays, immersive displays, virtual reality

- Multi-GPU systems such as PC clusters, NVIDIA QuadroPlex, etc.

GPU computing using custom module programmed using Avizo/Amira XPand and GPU API

- GPU clock frequency, number of GPU cores (stream processors)
- Multi-GPU systems such as NVIDIA Tesla
- CUDA support

©2016. We are constantly improving the performance of our products—all specifications are subject to change without notice. FEI, the FEI logo, Amira, Avizo, and Open Inventor are trademarks of FEI Company or its affiliates. All other trademarks belong to their respective owners. FL0034-03-2016



LANIKA SOLUTIONS PRIVATE LIMITED

TF-04, Gold Signature, No. 95, Mosque Road, Frazer Town, Bangalore - 560 005, INDIA

Phone: +91 – 80 – 2548 4844 Fax: +91 – 80 – 2548 4846 Email: info@lanikasolutions.com www.lanikasolutions.com