# Hardware Optimization for Immersive Simulation & Photogrammetric Environment Generation

**Jonathan Hawes, Karl Rosenberger**
**RAVE Computer**
**Sterling Heights, MI**
**jhawes@rave.com, krosenberger@rave.com**

## ABSTRACT

Civilian and Department of Defense (DoD) organizations are leveraging immersive simulation to create a digital twin of the world. This effort to maximize training and capability invokes difficulties when used in the field due to the processing power necessary to reconstruct photogrammetric environments in short suspense. Edge processing typically overlooks two critical areas (1) the hardware best suited to accomplish the task and (2) the most advantageous strategy for integration between the processing system and the software which will power the experience. With optimization in these areas, significant decreases in processing time can be achieved while decreasing the size, weight, and power consumption (SWaP) of the system. This allows for man-portable systems which can be employed directly with the user for training and operational use. This paper will address the issues stated above by sourcing firsthand data through hardware and photogrammetric software processing benchmarks across three hardware topologies: single node, dual node, and quad node configurations. This data will then be aggregated into the developmental foundation for a new system topology which will offer server grade reconstruction speeds in a workstation form factor. This data will showcase the effects of hardware optimization through decreased processing times, weight, power consumption, and heat generation. A real-world scenario will illustrate the mission-critical system performance and its potential infusion into the One World Terrain ecosystem. Finally, the real-world benefit of focusing on these key components in simulation-based systems for training will be explained along with courses of action for implementation by warfighters at the edge. Converging the advancements in software architecture with the evolution of central processing units (CPU) and graphics processing units (GPU), predicted by Moore's Law, we will set the stage for the next epoch of photogrammetric systems and software.

## ABOUT THE AUTHORS

**Jonathan Hawes** is a retiring Special Forces Intelligence Sergeant with over 17 years in the United States Army. During the last four years Jonathan has worked with the USASOC Force Modernization Center (UFMC) and Synthetic Training Environment-Cross Functional Team (STE-CFT) as the Chief of Tactical Collection, Processing, and Integration for One World Terrain: concentrating on the tactical application and advancement of edge processing and visualization capabilities. Mr. Hawes has worked alongside the University of Southern California Institute for Creative Technologies on multiple advancements in simulation & modeling to support both operational and training scenarios for the Department of Defense.

**Karl Rosenberger**, Chief Technologist for RAVE Computer, is a leader in innovation with over 35 years of experience designing purpose-built computers. His innovative system designs are currently in use by the Army Futures Command Synthetic Training Environment for One World Terrain photogrammetry processing. Over the last two decades Karl has forged strong partnerships with industry leading technology providers including Intel, NVIDIA AMD, Microsoft and RedHat. Mr. Rosenberger's revolutionary designs have been showcased at I/ITSEC for nearly 20 years.

# Hardware Optimization for Immersive Simulation & Photogrammetric Environment Generation

**Jonathan Hawes, Karl Rosenberger**
**RAVE Computer**
**Sterling Heights, MI**
**jhawes@rave.com, krosenberger@rave.com**

## PHOTOGRAMMETRY AT THE POINT OF NEED

### From Data Center to Epicenter

Just a few decades separated the development of the world's first photograph and the inception of an idea that would fundamentally change the geospatial world. In 1867, 155 years ago, German geographer Otto Kersten and German engineer Albrecht Meydenbauer conceptualized the idea of using photography to determine the distances between various objects in both mapping and surveying (Grimm, 1980). Since then, technology has evolved at an unprecedented rate. From cellphones and portable light detection & radar (LiDAR) sensors to aircraft and satellite collection methods, photography and data collection has become a part of everyday life. The increased capabilities of small form factor unmanned aerial systems (UAS) have made photogrammetry and remote sensing accessible to the masses. As collection methods have evolved, so to have the ways in which we process data.

Over the last 30 years, the amount of processing that can be done outside a data center has followed an exponential growth trend thanks to advances in technology. In 1993 Intel shocked the world with the release of its Pentium processor which brought a clock speed of 66MHz to the masses (Tom's Hardware, 2006). Just over a decade later Advanced Micro Devices (AMD) released the Duron 1800 which raced past other CPU offerings with a 1.8GHz base clock. By 2010 CPU manufactures had already surpassed the dual core and even quad core capabilities. Processing power was hitting unprecedented heights, and nothing encompasses this better than the Intel Core i7 980X Extreme edition and AMD Phenom II X6 1100T Black Edition. These titans boasted a staggering 6c/12t 3.33GHz base clock and 6c/6t 3.3GHz base clock respectively. For the average consumer these types of processing capabilities seemed unnecessary, but for industries like remote sensing and photogrammetry these numbers were not only palpable but drove software developers to create heavier and more hardware intensive offerings.

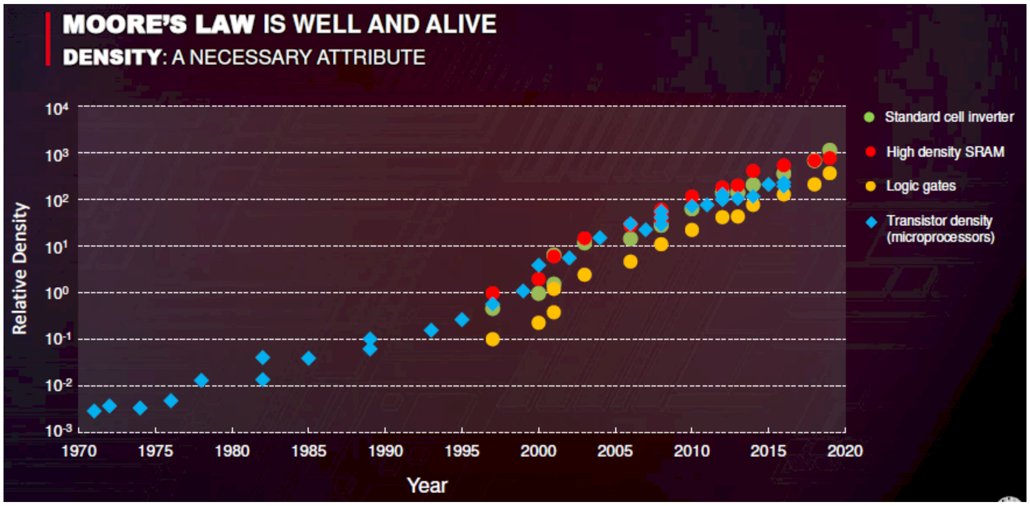


Figure 1 Moore's Law from 1970-2020. (Image taken from The Next Platform)

By 2020 the number of cores, threads, and base clock speeds had reached a level that dwarfed their predecessors. Intel had hit the market with their Xeon Phi 7295 in 2017 touting 72c/288t at 1.5GHz (Intel Corporation, 2022). AMD followed this flagrant display of technological prowess by releasing their flagship Ryzen Threadripper 3990X with an outrageous 64c/128t running at an astounding 2.9GHz base clock (Advanced Micro Devices Inc., 2022). These advancements have shown that Moore's law has held, relatively, for over 50 years. Between CPU progression and the ever-expanding capabilities of graphics processing units (GPU) work that was once destined for stacks of servers in a climate-controlled room is now being completed by first responders looking for survivors after a natural disaster, service members creating 3D reconstructions on the front lines before an operation, or even a farmer flying a drone to track crop grown over time. Investigators have taken to using photogrammetry to reconstruct vehicle accident scenes to determine accountability and fault. Researchers have taken to using LiDAR and photogrammetry to reconstruct ancient cities and historical sites to give the world a better idea of where we came from and how we got here. This capability has only recently been possible thanks to the continued development of microprocessors.

**Reconstructing the Real World One Drone at a Time**

As the ability to process large amounts of 3D data has expanded so too has the body of use in both civilian and military sectors. Companies like Maxar can produce sub-meter three dimensional offerings of nearly any place on earth. This has expanded the capabilities of first responders and Department of Defense entities as they plan and engage all over the planet. For those looking for a higher resolution reconstructions of specific smaller areas drones have become the flavor of choice. With the ability to capture a few hundred square meters up to tens of square kilometers enthusiasts and professionals alike have engaged with an industry that has exploded in growth in the last decade. From AeroVironment and FLIR creating state of the art military specific offerings to Parrot, DJI, SenseFly and others bridging the gap for consumers. The ability to collect, process, and disseminate information both in the field and at the office without the need for extensive computing environments has made it possible for the remote sensing industry to see continued growth. "The global remote sensing technology market size was valued at USD 12.40 billion in 2019 and is expected to grow at a compound annual growth rate (CAGR) of 11.6% from 2020 to 2027. Remote sensing technology is used for acquiring information about the earth's surface and to analyze its physical characteristics. It uses reflected and emitted light from aircraft and satellites without the need for physical contact with the surface area under study." (Grand View Research, 2020).

The process of creating a virtual representation of the natural world is complicated, consisting of mathematical algorithms, hundreds of thousands of lines of code, and countless hours of development by software engineers. Fundamentally however it is broken down into three phases. At the core of stereo-photogrammetry is the collection of the area or object. A sensor is used to collect a series of photos or video encompassing that which is to be reconstructed. Next a program conducts an Aerotriangulation (AT) of the photos. This process determines the intersection of points in space based on where they overlap within a series of photos. It is for this reason that within photogrammetry the more overlap there is between photos the more accurate the AT and overall reconstruction will be. Typically for the best results 60%-80% overlap is recommended. The final step in the process is the actual reconstruction. It is at this point that the program manipulates the photos and reconstructs them into the 3D model of the area or object that was originally captured. The result is a model of varying resolution depending on the distance between the sensor and the target along with the pixel density of the photos.

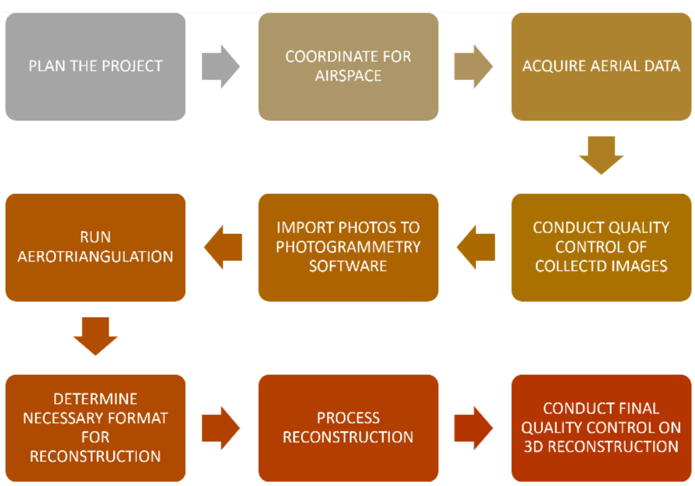

Figure 2 The Photogrammetric Process

This growth in the collection and dissemination market for remote sensing and photogrammetry shows there is no current trend marking a decline in the foreseeable future. Growth in collection and processing (speaking in terms of hardware) does have one fundamental issue. The current ability to process the collected data at the edge (specific to software) has not shown the same growth rate as the remote sensing industry. There are numerous programs commercially available that allow for static images, full-motion video, and LiDAR to be reconstructed into a 3D representation of the area they were collected from. The time it takes to process and reconstruct that data in some cases can take days or even weeks with current on-premises mobile and workstation hardware solutions. The current fix to this comes in two forms, on site servers allowing the data to be processed across multiple systems, or cloud enabled software offerings which are simply the use of someone else's server. In many cases this form of processing is acceptable, but it detracts from the accomplishments that have been made with technology by removing the possibility of edge computing. If we look specifically at the military use case for photogrammetry, the option of having multiple desktop computers at the front line to process 3D data rapidly is a non-sequitur when we think about how military operations are conducted. This issue of how to allow for rapidly reconstructed 3D models on the smallest platform possible is what sets the premise for the research which will be presented as the foundation for this paper. Up to this point the method for employing photogrammetric capabilities at the tactical or operational level has been to use commercially available software from companies like Bentley, Agisoft, or Pix4D and deploy them on commercial off the shelf mobile systems that are primarily used for gaming, and typically provide the computational hardware necessary to accomplish the mission, though at exceptionally long wait times.

## THE NEED FOR A PROCESSING REVOLUTION

### Defining a Path for Edge Photogrammetry

Though it is difficult to make significant changes to the hardware that is available in something mobile like a laptop, there is room for significant improvement in the desktop computing space. Determining the best configuration for processing 3D data is at the core of this research. There are dozens of different processors available on the market today, all sharing a singular theme. The more cores a processor has the lower the clock speed per core, with the reverse holding true, the lower the core count the higher the clock speed. This is mostly due to the heat produced by the processor based on the increased power consumption at higher core counts. This issue is what presents the initial case for this study; determining what the best configuration is for processing data, be it higher core counts across a single processor or lower core counts among a series of processors on separate systems. To determine the answer for this an experiment was conducted using three separate system topologies. The intent was to determine if there was a significant difference between the different topologies so that a determination could be made on how best to develop a system which supports rapid processing without the need for access to a distant server.

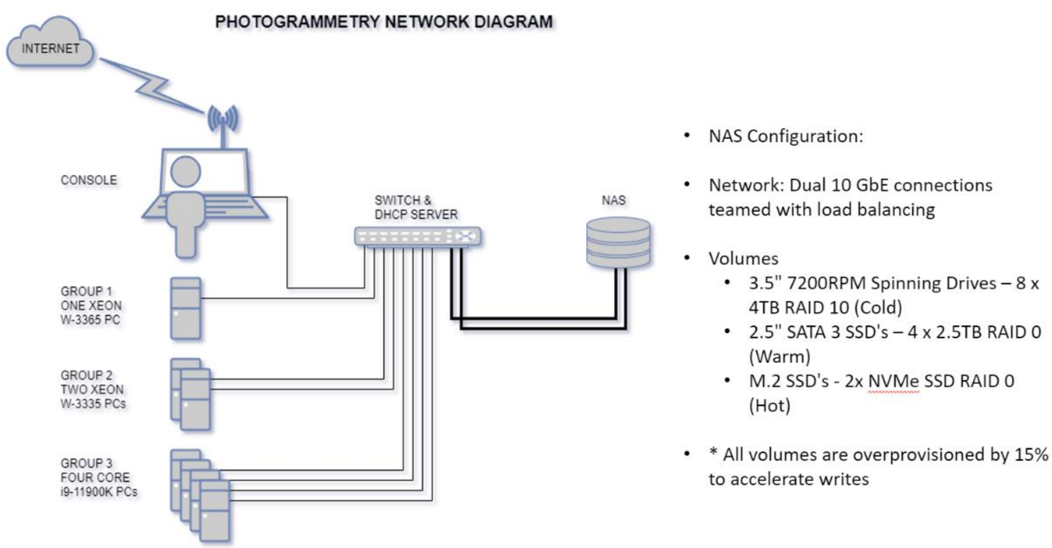


Figure 3 Photogrammetry Experiment Topology

The restrictions for the three topologies settled on 32 CPU processing cores, 64GB of DDR4 Random Access Memory (RAM), 1TB M.2 NVMe onboard storage, and 10Gbps access to the network attached storage which housed the data used for the reconstruction. Due to the limited availability of GPU's a single NVIDIA RTX A5000 was used on each system. After speaking with subject matter experts in the remote sensing and photogrammetric industry the consensus was that the GPU layout would not have a significant impact on the collected data compared to the cost of additional GPUs per system. To maintain the highest clock speeds per CPU core, an Intel Xeon W-3365 32c/64t at 2.7GHz CPU was used for the single system topology, the Intel Xeon W-3335 16c/32t at 3.4GHz was used for the dual system topology, and the Intel Core i9 11900K 8c/16t at 3.5GHz was used for the quad system topology. RAM was installed at 64GB of DDR4 across all systems. The power draw for each system was noted for each overall topology at 999 watts/8.3 amps, 1538 watts/12.8 amps, and 2776 watts/23.1 amps. All testing was conducted in a controlled environment with processors being water cooled and remaining system components relying on air cooling. Initial benchmarks were run across all systems using commercially available software to determine if there were any issues with the selected hardware or any outliers or exceptional differences in the results. Of the nine tests that were run the only excessive difference that was recorded was with the 11900K system, which far outperformed the other two configurations (during single system testing) in the Unigine Superposition 1080p Extreme Open GL test. The test was performed an additional three times to ensure accuracy showing similar results each time.

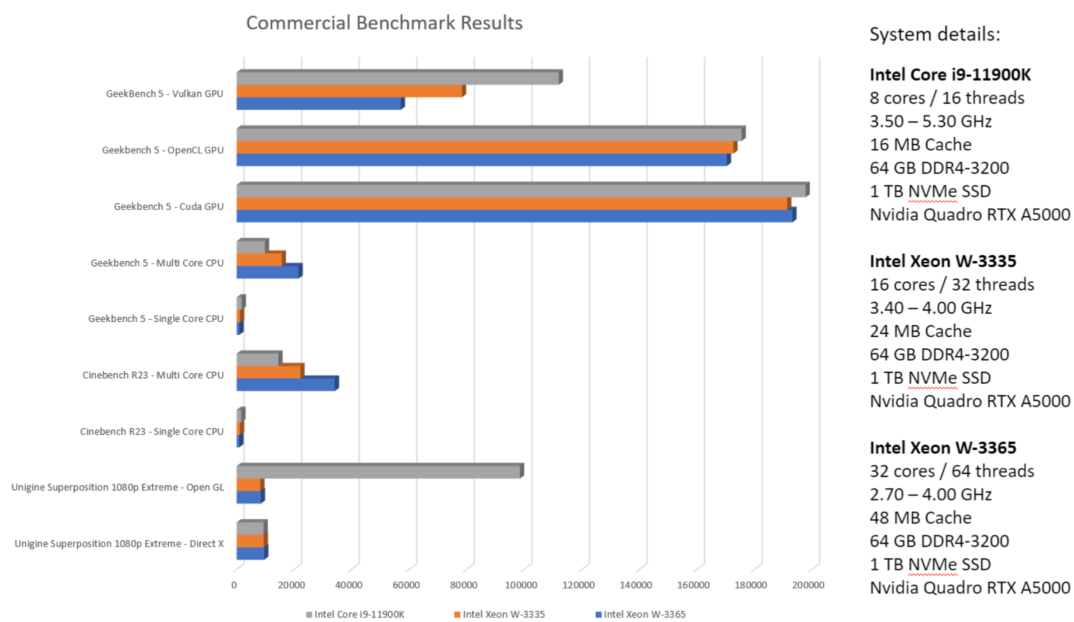


Figure 4 Commercial Benchmark Results

Once the benchmarks were completed and all hardware was found to be in working order a software platform was chosen with which to process the test data. As Bentley's Context Capture Center software is widely regarded as a leader in the photogrammetric space and allows for both single and multi-node processing it was selected for this experiment. The data set used consisted of 1,154 photos covering an area of 372,724m$^2$ which were collected by a DJI Phantom 3 UAS with a FC6310 13.2mm sensor that provided a 19mm/pixel average ground resolution. The AT across all topologies was completed using default settings to represent a standard that could be regarded as the common use case for an end user. Following the AT, the reconstructions were performed with regular planar grid tiling set to 16GB of RAM per tile which is consistent with the recommended usage set by Bentley for their software. The output format was set as 3MX which allows for high fidelity visualization within the software both during the process and after the reconstruction has finished. Following the completion of the AT and 3MX reconstructions all pertinent data including processing time, resolution, and hardware usage were recorded.

| | SINGLE NODE | DUAL NODE | QUAD NODE |
|---|---|---|---|
| **AEROTRAINGULATION PROCESSING TIME** | 00 DAYS<br>00 HOURS<br>40 MINUTES<br>27 SECONDS | 00 DAYS<br>00 HOURS<br>28 MINUTES<br>08 SECONDS | 00 DAYS<br>00 HOURS<br>25 MINUTES<br>34 SECONDS |
| **3MX PROCESSING TIME** | 00 DAYS<br>19 HOURS<br>11 MINUTES | 00 DAYS<br>05 HOURS<br>32 MINUTES | 00 DAYS<br>02 HOURS<br>41 MINUTES |

Figure 5 Photogrammetry AT/3MX Processing Results

## FINDINGS & CONCLUSION

### What's Good for the Many, Is not for the Few

The results of the testing were very telling of the way in which 3D data is processed. All three methods produced a final reconstruction that provided a sub-centimeter 3D model with incredibly high fidelity. The issue however stems from the need for multiple systems to produce the data rapidly. The dataset that was used for this experiment is relatively small in comparison to what can and has been collated for use by service members in the field. The question then is how to provide a portable solution which can meet or exceed the processing times of the 4-system topology which was the clear victor in this experiment. There will always be a use case for laptop systems to provide the capability for the man or women in the foxhole who has zero reach back to a rear echelon, but the goal is to determine a way for a service member to collect an area prior to launching a mission and provide them near real time data with which they can execute their mission. There are two schools of thought on this; the first being to consolidate four systems into a very small form factor which can fit into a vehicle or be man portable. The second is to virtualize the four systems into a small form factor which can be employed in the same way. Both paths have merit but contain their own sets of issues and obstacles ranging from heat generation and dissipation to ruggedization and deployment concerns in the extreme environments our women and men in the Armed Forces frequently find themselves. In addition to this, the larger the initial data set is the more compute intensive the overall reconstruction will be, demanding more power, more time, and generating more heat.

The way forward, now that this information has been collected, is to determine a feasible way to consolidate the incredible amount of processing power necessary for rapid reconstruction of 3D data which can then be implemented into current and future visualization systems for the Department of Defense and other entities such as the Department of Homeland Security, FEMA, and others. To do this, a succinct and phased process of continued development is necessary. Through a joint effort with partners such as Bentley Systems, NVIDIA, and the United States Army Synthetic Training Environment, the next step is to finalize the development of both physical and virtualized small form factor systems which mimic a quad node topology yet account for the issues stated within this paper. Along with this is the optimization of software and ensuring rapid and seamless integration of both the processing system and the generated models into the existing architecture of STE-IS, and those systems used by agencies like the National Geospatial Intelligence Agency and Army Geospatial Center. In doing this it will establish a precedent for the rapid integration of this evolved capability so it can be fielded to the force and increase both our training & readiness along with the operational lethality of the United States.

## ACKNOWLEDGEMENTS

RAVE would like to express its sincere gratitude to Bentley Systems and NVIDIA for their collaboration on this paper and the experimentation that led to the findings it contains. Over the last few years, the need for remote sensing and photogrammetric capabilities has increased dramatically. With this comes the need for revolutionary ideas and the painstaking process of taking theory to reality. We are excited for this initial phase experimentation and look forward to continuing to work with our partners as we move into the second and third phase of this process.

## REFERENCES

Advanced Micro Devices Inc. (2022, 04 28). *Processor Specifications*. Retrieved from AMD: https://www.amd.com/en/products/specifications/processors

Banfi, F., & Previtali, M. (2021). Human-Computer Interaction Based on Scan-to-BIM Models, Digital Photogrammetry, Visual Programming Language and eXtended Reality (XR). *Applied Sciences*.

Barrile, V., Bilotta, G., & Pozzoli, A. (2016). 3D Modeling: Comparison of some Innovative3 Techniques of PHotogrammetry. *WSEAS Transactions on Signal Processing*, 281-291.

Chen, M., Feng, A., Hou, Y., McCullough, K., Prasad, P., & Soibelman, L. (2021). Ground Material Segmentation for UAV-based Photogrammetric 3D Data: A 2D-3D Hybrid Approach. *Interservice/Industry Training, Simulation, and Education Conference*, 1-12.

Chen, M., Feng, A., McCullough, K., Prasad, P., Mcalinden, R., & Soibelman, L. (2020). 3D Photogrammetry Point Cloud Segmentation Using a Model Ensembling Framework. *Journal of Computing in Civil Engineering*, 1-20.

Feldman, M. (2019, September 13). *TSMC Thinks It Can Uphold Moore's Law For Decades*. Retrieved from The Next Platform: https://www.nextplatform.com/2019/09/13/tsmc-thinks-it-can-uphold-moores-law-for-decades/

Grand View Research. (2020, September 01). *Remote Sensing Technology Market Size, Share & Trends Analysis Report By Technology (Active, Pasive), By Application (Agricutlure, Military, Disaster Management, Weather), By Platform, By Region, And Segment Forecasts, 2020-2027.* San Francisco: Grand View Research. Retrieved from Grand View Research: https://www.grandviewresearch.com/industry-analysis/remote-sensing-technologies-market

Grimm, K. A. (1980). The Origin of the Term Photogrammetry. *Hamburg ISPRS-Congress*, 53-60. Retrieved from https://www.isprs.org/society/history/Grimm-The-Origin-of-the-Term-Photogrammetry.pdf

Guidi, G., Gonizzi, S., & Micoli, L. (2014). Image Pre-Processing for Optimizing Automated Photogrammetry Performances. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 145-152.

Howland, M., Kuester, F., & Levy, T. (2014). Photogrammetry in the Field: Documenting, Recording, and Presenting Archaeology. *Mediterranean Archaeology and Archaeometry*, 101-108.

Intel Corporation. (2022, 04 28). *Intel Processors*. Retrieved from Intel: https://www.intel.com/content/www/us/en/products/details/processors.html

Tom's Hardware. (2006, January 14). *Intel drops "Pentium" Brand*. Retrieved from tom'sHARDWARE: https://www.tomshardware.com/reviews/intel-drops-pentium-brand,1832-2.html