GETTING IN THE WAY : THE ARCHITECTURE OF ADDRESS pdf

1: Getting into Architecture? | Forum | Archinect

Information about planning a visit to the Getty Center and the Getty Villa in Los Angeles. Admission is always free.

While it seems like a drag working 8: This was the case this past Friday when I found myself alone in the office for several hours. Most of the time I use the peace and quiet to minimize the work that I would have to do over the weekend â€¦ in this case, I decided to take some photos so I could write this very particular blog post. Our office has grown to 11 people and we are at max capacity at the moment. I was up at the office working on some billing, the most exciting sort of work to stay late for, and I started thinking about how the personality of our office is really starting to round itself into proper form. Maybe my contribution is simply getting out of the way. I sit in the far back corner of the office. I have plenty of room and seem to expand out in all directions to fill it up. A few of us has made this memento an image magnet, on the far right, you can see my additions from my most recent trip to the Stiftung Bauhaus Dessau, and the image of Mt. Desert Island in Maine from this summer. This is a sketch from my desk â€” generated a few weeks ago for a fairly significant addition renovation project we have underway to a Robert A. Stern house here in Dallas. Whenever I look back to this wall prior to the addition of this shelf, I have a hard time wondering why it took so long to put it into place. One of the projects that seems to absorb the most attention in the office, is the addition and renovation of a 1. Everybody is getting in on the Post-It note action. We also have loads of as-built drawings. This particular associate has a very specific set of skills â€¦ being thorough and typically being right. I mentioned that we had a film crew in the office last week, although that might not actually be correct as it was one person can 1 person be a crew? Any guesses as to the things that received the most attention in the office?!? It was the models â€¦ but of course, you knew that.

GETTING IN THE WAY : THE ARCHITECTURE OF ADDRESS pdf

2: Featured Blogs | Archinect

The building at the top of the stairs is the Entrance Hall. The sculpture on the steps is called Air, designed by artist Aristide Maillol. The fact is that the Getty Center's architect Richard Meier did such an outstanding job of creating a public space that people get surprised.

Oct 19, 11 You are gonna work more than any other people you can imagine all the time until you become 50 years old. They are gonna say something positive thing about your work only in front of you if you are famous though or after you died. How can you handle this? You must love the architecture more than you love your first love, your girl friend, your dad, your brother, sister, mom, god or anything you can think of. You need to be a heartless fuck to be tougher or stronger than those people. Try to go to the architecture school first. If you cannot enjoy it and love it so much until 3rd year, you should get the fuck out of there ASAP because it will not be changed in the professional field. Oct 20, 11 4: Ask yourself, is this the only thing I can do as a career and be passionate about? Oct 20, 11 6: There are many other creative professions out there that take significantly less school, allow you time to have a social life and are, generally speaking, far more rewarding. Oct 20, 11 8: Honestly, the field of architecture is incredibly diverse. Architects are sort of on a sliding scale between an artist and an engineer. Some of us are incredibly artistic and talented, but lack technical abilities. Others are the opposite. And some consider themselves a jack of all trades. You really can have success anywhere in between. Salary is just fine. Hopefully the current doom and gloom will pass once you graduate school. How can a person become successful? Do what you want to do. Do what interests you. You will put in twice the effort and succeed twice as much if you pursue these things. Take art classes if you can. Anything, drawing, painting, photography. What information must you know? I would just try and push yourself in whatever direction interests you. There are a lot of different ways to practice architecture. I work in a planning and project management department at a University. I have a wife and two young sons, and the easier work load allows me to take part in their lives like I want to. I look forward one day to having my own firm, but this is a great intermediate place to be. All that to say, you can make it what you want. But to get into school, they all want to see a good portfolio. Any kind of art work will do, so long as you show you are a creative person. God bless your efforts Oct 20, 11 I would suggest you try to get a summer intern position in a local firm to see the profession from the inside and make an informed decision on your higher education major. If history is any indicator, the economy will be puttering along for another years. Oct 20, 11 1: See quote and links from a recent Compensation Survey below. Of course, many are leaving the field permanently. And some of the lost jobs were other staff. Still, the extent of industry contraction over such a short duration is difficult to fathom. For this reason, the competition will remain insane for quite a while. It will also likely continue to stagnate wages and benefits because there is much more supply than demand. By the end of , firm employment levels fell to under ,, according to the U. Department of Labor, producing more than a 25 percent decline in payroll positions over this period. What fields of engineering involving buildings and structures currently has the most employment rate or lowest unemployment rate. And what is its salary. Oct 20, 11 5: CNN lists it in as one of the top jobs. I know computer engineers are in high demand though and my research tells me that most people who end up in the field never even studied programming in HS based on one site. So computer engineer and relate jobs are on ym list. Civil engineering and other types are too. But I am looking at performing Energy Star Audits for some commercial properties: Oct 20, 11 9:

GETTING IN THE WAY : THE ARCHITECTURE OF ADDRESS pdf

3: Visit the Getty

We are in a "new normal" (as the wise folk at Pimco famously proclaimed), meaning that no one really knows if/when things will go back to the way it goes. Look above at the numbers schools are graduating.

Note A failed NAT gateway is automatically deleted after a short period; usually about an hour. The following table lists the possible causes of the failure as indicated in the Amazon VPC console. Displayed error Remedial steps Subnet has insufficient free addresses to create this NAT gateway The subnet you specified does not have any free private IP addresses. You can check how many IP addresses are available in your subnet by going to the Subnets page in the Amazon VPC console, and viewing the Available IPs box in the details pane for your subnet. To create free IP addresses in your subnet, you can delete unused network interfaces, or terminate instances that you do not require. Create and attach an internet gateway to your VPC. For more information, see Creating and Attaching an Internet Gateway. Elastic IP address eipalloc-xxxxxxxx is already associated The Elastic IP address that you specified is already associated with another resource, and cannot be associated with the NAT gateway. If you do not require the Elastic IP address for that resource, you can disassociate it. Alternatively, allocate a new Elastic IP address to your account. Network interface eni-xxxxxxxx, created and used internally by this NAT gateway is in an invalid state. There was a problem creating or using the network interface for the NAT gateway. You cannot fix this error. Try creating a NAT gateway again. Check the status of your NAT gateway. A status of Pending, Available, or Deleting counts against your limit. We cannot support NAT gateways in these zones. You can create a NAT gateway in another Availability Zone and use it for private subnets in the constrained zone. You can also move your resources to an unconstrained Availability Zone so that your resources and your NAT gateway are in the same Availability Zone. A NAT gateway only passes traffic from an instance in a private subnet to the internet. Instances in Private Subnet Cannot Access Internet If you followed the preceding steps to test your NAT gateway and the ping command fails, or your instances cannot access the internet, check the following information: Check that the NAT gateway is in the Available state. If the NAT gateway is in a failed state, there may have been an error when it was created. The NAT gateway must be in a public subnet with a route table that routes internet traffic to an internet gateway. For more information, see Creating a Custom Route Table. Your instance must be in a private subnet with a route table that routes internet traffic to the NAT gateway. For more information, see Updating Your Route Table. Check that there are no other route table entries that route all or part of the internet traffic to another device instead of the NAT gateway. Ensure that your security group rules for your private instance allow outbound internet traffic. For the ping command to work, the rules must also allow outbound ICMP traffic. Note The NAT gateway itself allows all outbound traffic and traffic received in response to an outbound request it is therefore stateful. Ensure that the network ACLs that are associated with the private subnet and public subnets do not have rules that block inbound or outbound internet traffic. For the ping command to work, the rules must also allow inbound and outbound ICMP traffic. Note You can enable flow logs to help you diagnose dropped connections because of network ACL or security group rules. If you are using the ping command, ensure that you are pinging a website that has ICMP enabled. If not, you will not receive reply packets. To test this, perform the same ping command from the command line terminal on your own computer. Check that your instance is able to ping other resources, for example, other instances in the private subnet assuming that security group rules allow this. To check if the endpoint is sending fragmented TCP packets, use an instance in a public subnet with a public IP address to do the following: Trigger a response large enough to cause fragmentation from the specific endpoint. Use the tcpdump utility to verify that the endpoint is sending fragmented packets. Important You must use an instance in a public subnet to perform these checks; you cannot use the instance from which the original connection was failing, or an instance in a private subnet behind a NAT gateway or a NAT instance. Diagnostic tools that send or receive large ICMP packets will report packet loss. For example, the command ping -s example. In this case, your instance is accessing the internet using a different device, such as an internet gateway. In the route table of the subnet in which your instance is located, check the following

information: Ensure that there is a route that sends internet traffic to the NAT gateway. To prevent the connection from being dropped, you can initiate more traffic over the connection or enable TCP keepalive on the instance with a value smaller than seconds. Cannot Initiate More Connections to a Destination You might have reached the limit for simultaneous connections. If your instances in the private subnet create a large number of connections, you may reach this limit. You can do one of the following: Create additional NAT gateways in the public subnet and split your clients into multiple private subnets, each with a route to a different NAT gateway. Limit the number of connections your clients can create to the destination. To release the capacity, close idle connections. We currently do not support resource-level permissions for any of the ec2: Cost allocation tags are supported for NAT gateways, therefore you can also use tags to organize your AWS bill and reflect your own cost structure.

GETTING IN THE WAY : THE ARCHITECTURE OF ADDRESS pdf

4: Architecture Studio Vignettes | Life of an Architect

The most successful architecture firms have multiple 'client-getting' channels. The key to successful marketing for architects is to have several of these channels working for your business. The more you use, the more you can be pro-active about which projects you take, instead of being reactive to the projects that come through the door.

History of Architecture Throughout ancient and medieval history, most of the architectural design and construction was carried out by artisans â€"such as stone masons and carpenters , rising to the role of master builder. Until modern times, there was no clear distinction between architect and engineer. In Europe, the titles architect and engineer were primarily geographical variations that referred to the same person, often used interchangeably. Paper was not used in Europe for drawing until the 15th century but became increasingly available after Pencils were used more often for drawing by The availability of both allowed pre-construction drawings to be made by professionals. Until the 18th-century, buildings continued to be designed and set out by craftsmen with the exception of high-status projects. Such licensure usually requires an accredited university degree, successful completion of exams, and a training period. The use of terms and titles and the representation of oneself as an architect is restricted to licensed individuals by law, although in general, derivatives such as architectural designer are often not legally protected. To practice architecture implies the ability to practice independently of supervision. The term building design professional or Design professional , by contrast, is a much broader term that includes professionals who practice independently under an alternate profession, such as engineering professionals, or those who assist in the practice architecture under the supervision of a licensed architect, such as architectural technologists and intern architects. In many places, independent, non-licensed individuals may perform design services outside the professional restrictions, such design houses and other smaller structures. Practice[edit] In the architectural profession, technical and environmental knowledge, design and construction management, and an understanding of business are as important as design. However, the design is the driving force throughout the project and beyond. An architect accepts a commission from a client. The commission might involve preparing feasibility reports, building audits, the design of a building or of several buildings, structures, and the spaces among them. The architect participates in developing the requirements the client wants in the building. Throughout the project planning to occupancy , the architect co-ordinates a design team. Structural , mechanical , and electrical engineers and other specialists, are hired by the client or the architect, who must ensure that the work is co-ordinated to construct the design. Design role[edit] The architect hired by a client is responsible for creating a design concept that meets the requirements of that client and provides a facility suitable to the required use. In that, the architect must meet with and question the client to ascertain all the requirements and nuances of the planned project. Often the full brief is not entirely clear at the beginning, entailing a degree of risk in the design undertaking. The architect may make early proposals to the client which may rework the terms of the brief. The program or brief is essential to producing a project that meets all the needs of the owner â€" it is a guide for the architect in creating the design concept. It is generally expected that the design proposal s is both imaginative as well as pragmatic, but the precise extent and nature of these expectations will vary, depending on the place, time, finance, culture, and available crafts and technology in which the design takes place. Designing buildings is a very complex and demanding undertaking, no matter what the scale of the project might be. A strong degree of foresight is a prerequisite. Any design concept must at a very early stage in its generation take into account a great number of issues and variables which include qualities of space s , [8] the end-use and life-cycle of these proposed spaces, connections, relations, and aspects between spaces including how they are put together as well as the impact of proposals on the immediate and wider locality. Selection of appropriate materials and technology must be considered, tested and reviewed at an early stage in the design to ensure there are no setbacks such as higher-than-expected costs which may occur later. The site and its environs, as well as the culture and history of the place, will also influence the design. The design must also countenance increasing concerns with environmental sustainability. The architect may introduce intentionally or not , to greater or lesser degrees, aspects of mathematics and architecture , new or current architectural

theory , or references to architectural history. A key part of the design is that the architect often consults with engineers, surveyors and other specialists throughout the design, ensuring that aspects such as the structural supports and air conditioning elements are coordinated in the scheme as a whole. The control and planning of construction costs are also a part of these consultations. Coordination of the different aspects involves a high degree of specialized communication, including advanced computer technology such as BIM Building Information Management , CAD, and cloud-based technologies. At all times in the design, the architect reports back to the client who may have reservations or recommendations, introducing a further variable into the design. Architects deal with local and federal jurisdictions about regulations and building codes. The architect might need to comply with local planning and zoning laws, such as required setbacks, height limitations, parking requirements, transparency requirements windows , and land use. Some established jurisdictions require adherence to design and historic preservation guidelines. Health and safety risks form a vital part of the current design, and in many jurisdictions, design reports and records are required which include ongoing considerations such as materials and contaminants, waste management and recycling, traffic control and fire safety. Means of design[edit] Previously, architects employed drawings [6] to illustrate and generate design proposals. While conceptual sketches are still widely used by architects, [9] computer technology has now become the industry standard. Increasingly, computer software such as BIM is shaping how architects work. Renewable energy sources may be developed within the proposed building or via local or national renewable energy providers. As a result, the architect is required to remain abreast of current regulations which are continually tightening. Some new developments exhibit extremely low energy use. Construction role[edit] As the design becomes more advanced and detailed, specifications and detail designs are made of all the elements and components of the building. Techniques in the production of a building are continually advancing which places a demand on the architect to ensure that he or she remains up to date with these advances. Architects typically put projects to tender on behalf of their clients, advise on the award of the project to a general contractor , facilitate and then administer a contract of agreement which is often between the client and the contractor. Depending on the type of contract utilized, provisions for further sub-contract tenders may be required. The architect may require that some elements are covered by a warranty which specifies the expected life and other aspects of the material, product or work. In most jurisdictions, prior notification to the relevant local authority must be given before commencement on site, thus giving the local authority notice to carry out independent inspections. The architect will then review and inspect the progress of the work in coordination with the local authority. The architect will typically review contractor shop drawings and other submittals , prepare and issue site instructions, and provide Certificates for Payment to the contractor see also Design-bid-build which is based on the work done to date as well as any materials and other goods purchased or hired. In the United Kingdom and other countries, a quantity surveyor is often part of the team to provide cost consulting. With very large, complex projects, an independent construction manager is sometimes hired to assist in the design and to manage construction. In many jurisdictions, mandatory certification or assurance of the completed work or part of works is required. This demand for certification entails a high degree of risk - therefore, regular inspections of the work as it progresses on site is required to ensure that is in compliance with the design itself as well as with all relevant statutes and permissions. Alternate practice and specializations[edit] Recent decades have seen the rise of specializations within the profession. Many architects and architectural firms focus on certain project types for example, healthcare, retail, public housing, event management , technological expertise or project delivery methods. Some architects specialize as building code, building envelope , sustainable design , technical writing , historic preservation US or conservation UK , accessibility and other forms of specialist consultants. Many architects elect to move into real estate property development , corporate facilities planning, project management , construction management, interior design , city planning, or other related fields.

GETTING IN THE WAY : THE ARCHITECTURE OF ADDRESS pdf

5: How to Use Names, Titles, and Forms of Address

The only architecture magazine with a social conscience, your subscription is designed to deepen your knowledge and challenge you with reflections on the social purpose of architecture. Subscribe now 10+ subscriptions.

For example, the linear address h can have the segmented addresses 06EFh: This could be confusing to programmers accustomed to unique addressing schemes, but it can also be used to advantage, for example when addressing multiple nested data structures. Therefore, real mode can just as well be imagined as having a variable length for each segment, in the range 1 to bytes, that is just not enforced by the CPU. The leading zeros of the linear address, segmented addresses, and the segment and offset fields are shown here for clarity. They are usually omitted. Both were packaged in pin DIP packages; even with only 20 address lines, the address and data buses were multiplexed to fit all the address and data lines within the limited pin count. A segment value of 0Ch 12 would give a linear address at C0h in the linear address space. The address offset can then be added to this number. Such address translations are carried out by the segmentation unit of the CPU. On the , these address accesses were wrapped around to the beginning of the address space such that In bit real mode, enabling applications to make use of multiple memory segments in order to access more memory than available in any one 64K-segment is quite complex, but was viewed as a necessary evil for all but the smallest tools which could do with less memory. The root of the problem is that no appropriate address-arithmetic instructions suitable for flat addressing of the entire memory range are available. The memory model concept derives from the setup of the segment registers. This section needs additional citations for verification. Please help improve this article by adding citations to reliable sources. Unsourced material may be challenged and removed. August Three segments in protected mode memory click on image to enlarge , with the local descriptor table. Instead, the bit segment registers now contain an index into a table of segment descriptors containing bit base addresses to which the offset is added. To support old software, the processor starts up in "real mode", a mode in which it uses the segmented addressing model of the There is a small difference though: With the addition of the HMA, the total address space is approximately 1. Moreover, it still necessitated dividing memory into 64k segments like was done in real mode. This limitation can be worked around on bit CPUs which permit the use of memory pointers greater than 64k in size, however as the Segment Limit field is only bit long, the maximum segment size that can be created is 16MB although paging can be used to allocate more memory, no individual segment may exceed 16MB. This method was commonly used on Windows 3. Thus, it was still necessary to place all code that performs API calls in 64k segments. Once protected mode is invoked, it could not be exited except by performing a hardware reset. The segment selector must be located in one of the segment registers. It then performs the privilege check: All privilege levels are integers in the range 0â€”3, where the lowest number corresponds to the highest privilege. If the inequality is false, the processor generates a general protection GP fault. Otherwise, address translation continues. The processor then takes the bit or bit offset and compares it against the segment limit specified in the segment descriptor. If it is larger, a GP fault is generated. Otherwise, the processor adds the bit segment base, specified in descriptor, to the offset, creating a linear physical address. The privilege check is done only when the segment register is loaded, because segment descriptors are cached in hidden parts of the segment registers. Also, importantly, address offsets are bit instead of bit , and the segment base in each segment descriptor is also bit instead of bit. The general operation of the segmentation unit is otherwise unchanged. The paging unit may be enabled or disabled; if disabled, operation is the same as on the If the paging unit is enabled, addresses in a segment are now virtual addresses, rather than physical addresses as they were on the That is, the segment starting address, the offset, and the final bit address the segmentation unit derived by adding the two are all virtual or logical addresses when the paging unit is enabled. When the segmentation unit generates and validates these bit virtual addresses, the enabled paging unit finally translates these virtual addresses into physical addresses. The physical addresses are bit on the , but can be larger on newer processors which support Physical Address Extension. Unlike protected mode, CPUs can be put back into real mode merely by clearing a bit in the CR0 control register. Later developments[edit] The x architecture does not use

segmentation in long mode bit mode. Four of the segment registers: The segment registers FS and GS can still have a nonzero base address. This allows operating systems to use these segments for special purposes. Unlike the global descriptor table mechanism used by legacy modes, the base address of these segments is stored in a model-specific register. The x architecture further provides the special SWAPGS instruction, which allows swapping the kernel mode and user mode base addresses. On x64, the CPU powers on into real mode and is indistinguishable from a bit Pentium 4. When long mode is operating, bit instructions and virtual x86 mode are disabled and protected mode disappears. Logical addresses can be explicitly specified in x86 assembly language , e. Most, but not all, instructions that use DS by default will accept an ES override prefix. Processor stack references, either implicitly e. For instance, the Linux kernel sets up only 4 general purpose segments:

GETTING IN THE WAY : THE ARCHITECTURE OF ADDRESS pdf

6: What is new in the SAP Community

Global Engagement. Travel the world in our Global Engagement Studios and Study Abroad Programs to experience cultures, climates, people, and places that will forever change the way you think about the critical roles of the future of architecture and design.

L2 cache is inclusive of L1 cache. It is okay to have a cache line in L2 only, but L1 cache lines must have a copy in L2. L2 cache associated with each core is of size KB. The cache is divided into two logical banks. L2 cache can deliver 64 bytes of read data to corresponding cores every two cycles and 64 bytes of write data every cycle. Multithreading The Intel Xeon Phi coprocessor uses time-multiplexed multithreading. The PPF thread picker determines the request to be sent to instruction cache and put into the prefetch buffer and the second one selects what entries to read from the prefetch buffer and send them to the decoder. It uses a smart round-robin mechanism to pick only threads that have work to do and avoids threads that are inactive due to various conditions like cache miss. One can control thread scheduling by putting delay loops in the thread that the thread picker will not schedule in the actual hardware which is useful for optimization.. In order to support multithreading, all architectural states are replicated four times. Micro architectural states of prefetch buffers, instruction pointers, segment descriptors, and exception logic are replicated four times as well. All memory stalls are converted into thread-specific flushes. Performance Considerations Run two or more threads to have the core completely busy. Integer operations and mask instructions are single cycle. Most vector instructions have four-cycle latency and single-cycle throughput, so having four threads one will not see the vector unit latency if all threads are executing vector instructions. Use four threads to hide vector unit latencies. When one writes to a GPR that is used as a base or index register, these instructions have to be spaced properly as the address generation is done in a separate stage in the pipeline. For example consider the following instruction sequence: Add rbx,4 Mov rax,[rbx] Here the calculation of the actual linear address from [rbx] is done in a separate stage before the memory fetch happens at line 2 above. In this case hardware will insert two clock delays between these two instructions. If running more than one thread, one may not see it as instructions from other threads can run during the dead clock cycles for another thread. The cores do not forward the store buffer contents, even if the data is available there requested by the load instructions. So a load followed by a store to the same location will need the store buffer to write to the cache and then readback from the cache. In this case the latency is four clock cycles. If there is a bank conflict, that is the U- and V-pipe try to access a memory location that resides on the same L2 cache bank, the core will introduce two clock-cycle delays for the V-pipe. When the cache line has to be replaced, it involves two clock-cycle delays. So if the code has a lot of vector instructions that miss cache, the core will have to put in a two cycle delay for each miss to do the replacement. Prefix Decode There are two classes of prefixes. The fast prefixes are decoded in 0 cycles with dedicated hardware blocks and include following prefixes: All other prefixes are decoded in two-cycle latencies and include 66 prefixes for selecting bit operand size, address size 67, lock prefix, segment prefix, and REP prefix for string instructions. Pairing Rules There are specific instructions that are pairable and thus can execute in both the U and V pipeline. Theseinstructions include independent single-cycle instructions. For pairing to occur, instructions cannot have both displacement and immediate, and the instructions cannot be longer than 8 bytes. Microcode instructions do not pair. Integer multiply operations are done in x87 floating point units; that is why it is long latency instructions 10 cycles in KNC. One needs to shift the integer operands to FP unitsto do the multiply and shift the results back to the integer pipeline. Probing The Core Measuring Peak Gflops In this section we will examine the core microarchitecture and pipelines described above through experiments. In these processors there are 60 cores available for computation while one of the cores services the operating system. The code is listed in the code listing 4. To start with the code needs to be designed such that the computation is core bound and not limited by the memory or cache bandwidth. To do so, we need the data to fit in the processor vector registers. This is limited to 32 for each thread in Intel Xeon Phi cores. Please see appendix A for a brief back ground on OpenMP. This file allows us to query various OpenMP parameters. Writing the code in OpenMP will allow us to run the code from one core to multiple

cores using an environment switch. Thus we can experiment with the effect of multiple threads in a core, to running the code on all cores to understand the compute power of the processor. Following the code, In line 42 and 43 we declare two constants. The first constant 16 indicates the number of double precisions we should be using for our computation. In this architecture the vector units can work on 8 DP number per cycle. As wehave seen in architecture description section, each of the four threads per core has a ready-to-run buffer prefetch buffer of two instructions deep as each core is able to issue two instructions per clock. Thus having 16 DP elements allow us to put a pair of independent instructions in the fill buffer. I have also made the number of iterations ITER multiple of , the number of maximum threads I shall be running to balance the workload. Line 52 declared three arrays that I shall be working on each in each thread. I have aligned the arrays to 64 byte cacheline boundaries so that they can be loaded efficiently by the compiler. If you look carefully, all threads will be writing back to array a[] in line 65 and 73 in the code listing below. This will cause a race condition and will not be useful in real application code. However, for illustrative purposes, it can be ignored at this time. In line 62 through 66 we warm up the cache and initialize openmp threads so that these overheads are not counted during the code timing loops that happen between lines 70 through For easy reference, if you look the code fragment below, line 70 is a repetition of the omp parallel for first encountered at line As such, all OpenMP overhead related to thread creation is captured at line 62 and the threads are reused at line The pragma vector aligned at line 72 tells the compiler that the arrays a, b and c are all aligned to 64 byte boundary for Intel Xeon Phi and does not need to do any special load manipulations needed for unaligned data. At line 78 and 79 the Gflops operations are computed. In addition, the Intel compiler allows us to control the code generation such that we can turn vectorization on or off. This help indicate to the array that the elements of the arrays a,b,c in this case are independent and thus allowing the compiler to vectorize the code which otherwise could be ambiguous to the compiler as for its vectorizability. To begin with we would need to build the code with vectorization turned off. To build with the Intel Compiler with vectorization turned off we will use the following command line: The source file, dpflops. Once the code is compiled, copy the file produces, dpflops. The operating system on the Intel Xeon Phi usually allocates some portion of the GDDR memory to be used as a ramdisk and hosts the file system on the card. You will also need to upload the dependent openmp library, libiomp5. Once the card is uploaded with the file, you can log onto the card using the command: However, as described in architecture section, the hardware cannot issue instructions back to back from the same thread in the core. To reach full execution unit utilization at least two threads must be running at all times. As the core still uses vector unit to perform scalar arithmetic, the code for scalar arithmetic is very inefficient. For each fma on a DP element, it has to broadcast the element to all the lanes. Operate on the vector register with a mask and then store the single element back to memory. We also notice that increasing threads per core does not improve performance as the instruction can be issued every cycle for this case. Now if we extend to threads utilizing all 60 cores we can achieve 86 Gflops as shown in the figure. If we recompile the code with following command line: We will get vectorized version of the code. In above generated compiler report you can see that line 65 and line 73 of the code listing 4. Figure 4 shows the results of experimentation with the vectorized double precision code. With 1 thread, we are able to 8. As described above this is due to the issue BW is not fully utilized thus execution unit is not utilized each cycle. In order to achieve so we need to use at least two threads. This kept execution unit fully utilized and was able to achieve near peak performance of We also observe in the figure that increasing number of threads to 3 and 4 did not improve performance since we already saturated the double precision execution unit on the core. Utilizing all 60 cores, with 2 threads each we can see that we could achieve Gflops which is near the theoretical peak of Gflops on this coprocessor. Since we are running threads with 2 threads per core on 60 cores, having affinity set to balanced, there will be two threads per core distributed across cores. We would expect the performance to be 2x the single precision. The changes to the code segment in code listing 1 will be to change double to float. However, we need to keep the float calculation same as the number of operations are the same. However, since the vector unit can work on 16 elements at a time with the float, in order to be able to fill up the dispatch queue for each thread, we would need twice as many elements as double, so I have adjusted the size to 32 elements instead of 16 elements as shown below. For single precision arithmetic, we did the same experimentation as before and the

output is captured below in Figure 5. However, adding more threads to core on this core bound code does not help. After changing the affinity to balanced, we allowed the threads to spread out to various cores first to make use of independent execution units. Making the number of threads 2xNumber of cores for this case , we were able to get the performance to GFlops. The lmbench was developed by the authors to make them portable and wholly written in C. That was the reason I picked the benchmark to explore the memory hierarchy of Intel Xeon Phi architecture.

GETTING IN THE WAY : THE ARCHITECTURE OF ADDRESS pdf

7: Evaluating a Software Architecture | Why Evaluate an Architecture? | InformIT

Read "So many ways of getting in the way: diversity in the molecular architecture of superantigen-dependent T-cell signaling complexes, Current Opinion in Immunology" on DeepDyve, the largest online rental service for scholarly research with thousands of academic publications available at your fingertips.

Marry your architecture in haste and you can repent in leisure. And Very Few Lead Bullets Either How can you be sure whether the architecture chosen for your software is the right one? The foundation for any software system is its architecture. No amount of tuning or clever implementation tricks will wring any of these qualities out of a poorly architected system. To put it bluntly, an architecture is a bet, a wager on the success of a system. Until recently, there were almost no methods of general utility to validate a software architecture. If performed at all, the approaches were spotty, ad hoc, and not repeatable. We can do better than that. This is a guidebook of software architecture evaluation. It is built around a suite of three methods, all developed at the Software Engineering Institute, that can be applied to any software-intensive system: Active Reviews for Intermediate Designs The methods as a group have a solid pedigree, having been applied for years on dozens of projects of all sizes and in a wide variety of domains. With these methods, the time has come to include software architecture evaluation as a standard step of any development paradigm. Evaluations represent a wise risk-mitigation effort and are relatively inexpensive. They pay for themselves in terms of costly errors and sleepless nights avoided. Whereas the previous chapter introduced the concept of software architecture, this chapter lays the conceptual groundwork for architectural evaluation. It defines what we mean by software architecture and explains the kinds of properties for which an architecture can and cannot be evaluated. The software architecture of a program or computing system is the structure or structures of the system, which comprise software components, the externally visible properties of those components, and the relationships among them. The architecture defines the components such as modules, objects, processes, subsystems, compilation units, and so forth and the relevant relations such as calls, sends-data-to, synchronizes-with, uses, depends-on, instantiates, and many more among them. The architecture is the result of early design decisions that are necessary before a group of people can collaboratively build a software system. One of the insights about architecture from Chapter 1 that you must fully embrace before you can understand architecture evaluation is this: Sooner or later everyone asks the question: Perhaps you asked it in one of the following ways: What is the difference between an architecture and a high-level design? Are details such as priorities of processes architectural? Why should implementation considerations such as buffer overflows be treated as architectural? Are interfaces to components part of the architecture? If I have class diagrams, do I need anything else? Is architecture concerned with run-time behavior or static structure? Is the operating system part of the architecture? Is the programming language? First, consider the definition of architecture that we quoted in Chapter 1 of this book. A software architecture concerns the gross organization of a system described in terms of its components, their externally visible properties, and the relationships among them. True enough, but it fails to explicitly address the notion of context. If the scope of my concern is confined to a subsystem within a system that is part of a system of systems, then what I consider to be architectural will be different than what the architect of the system of systems considers to be architectural. It has been said that algorithms are not architectural; data structures are not architectural; details of data flow are not architectural. Well, again these statements are only partially true. Some properties of algorithms, such as their complexity, might have a dramatic effect on performance. Some properties of data structures, such as whether they need to support concurrent access, directly impact performance and reliability. Some of the details of data flow, such as how components depend on specific message types or which components are allowed access to which data types, impact modifiability and security, respectively. So is there a principle that we can use in determining what is architectural? Our criterion for something to be architectural is this: It must be a component, or a relationship between components, or a property of components or relationships that needs to be externally visible in order to reason about the ability of the system to meet its quality requirements or to support decomposition of the system into independently implementable pieces. Here are some corollaries

of this principle: Architecture describes what is in your system. Architecture describes the part that is in. An architecture is an abstract depiction of your system. The information in an architecture is the most abstract and yet meaningful depiction of that aspect of the system. Given your architectural specification, there should not be a need for a more abstract description. That is not to say that all aspects of architecture are abstract, nor is it to say that there is an abstraction threshold that needs to be exceeded before a piece of design information can be considered architectural. The architecture bridges the gap between requirements and the rest of the design. If you feel that some information is critical for reasoning about how your system will meet its requirements then it is architectural. You, as the architect, are the best judge. On the other hand, if you can eliminate some details and still compose a forceful argument through models, simulation, walk-throughs, and so on about how your architecture will satisfy key requirements then those details do not belong. However, if you put too much detail into your architecture then it might not satisfy the next principle. An architectural specification needs to be graspable. The whole point of a gross-level system depiction is that you can understand it and reason about it. Too much detail will defeat this purpose. An architecture is constraining. It imposes requirements on all lower-level design specifications. I like to distinguish between when a decision is made and when it is realized. For example, I might determine a process prioritization strategy, a component redundancy strategy, or aset of encapsulation rules when designing an architecture; but I might not actually make priority assignments, determine the algorithm for a redun-dant calculation, or specify the details of an interface until much later. To be architectural is to be the most abstract depiction of the system that enables reasoning about critical requirements and constrains all subsequent refinements. If it sounds like finding all those aspects of your system that are architectural is difficult, that is true. It is unlikely that you will discover everything that is rchitectural up front, nor should you try. The earlier you find a problem in a software project, the better off you are. The cost to fix an error found during requirements or early design phases is orders of magnitudes less to correct than the same error found during testing. Architecture is the product of the early design phase, and its effect on the system and the project is profound. An unsuitable architecture will precipitate disaster on a project. Performance goals will not be met. Security goals will fall by the wayside. The customer will grow impatient because the right functionality is not available, and the system is too hard to change to add it. Schedules and budgets will be blown out of the water as the team scrambles to back-fit and hack their way through the problems. Months or years later, changes that could have been anticipated and planned for will be rejected because they are too costly. Plagues and pestilence cannot be too far behind. Architecture also determines the structure of the project: If it changes midstream because of some deficiency discovered late, the entire project can be thrown into chaos. It is much better to change the architecture before it has been frozen into existence by the establishment of downstream artifacts based on it. Architecture evaluation is a cheap way to avoid disaster. The methods in this book are meant to be applied while the architecture is a paper specification of course, they can be applied later as well , and so they involve running a series of simple thought experiments. They each require assembling relevant stakeholders for a structured session of brainstorming, presentation, and analysis. All told, the average architecture evaluation adds no more than a few days to the project schedule.

8: NAT Gateways - Amazon Virtual Private Cloud

So many ways of getting in the way: diversity in the molecular architecture of superantigen-dependent T-cell signaling complexes Author links open overlay panel Eric J Sundberg Yili Li Roy A Mariuzza Show more.

9: x86 memory segmentation - Wikipedia

The source IP address should be the Elastic IP address of your NAT gateway. You can get the Elastic IP address and private IP address of your NAT gateway by viewing its information on the NAT Gateways page in the Amazon VPC console.

GETTING IN THE WAY : THE ARCHITECTURE OF ADDRESS pdf

Biology of belief book Birmingham (1952 to 1968) Balanced living on a tightrope Subjects, citizens, and refugees Network analysis gate questions with solutions Library of American literature from earliest settlement to the present time Sight for books Migration Considerations for Cics Using Cics Cm, Cics Pa, and Cics Ia Gods, goddesses, and images of God in ancient Israel Camera shot list template The hidden alcoholic in general practice What is a buffer? Testing vs. your child. The witchs magic cloth. Social movements in twentieth-century Iran Reading selection : School, part 1 Passion And Value Of Salvation Decision-making for nurse educators Reflections in a mission mirror Japanese exotic cars Yvan Goll; an iconographical study of his poetry. Towards the Muslim apocalypse Thought of the Evangelical Leaders Out of the Italian Night The Ontario township Rock Power Guitar Pak (Rock Power Series) German philosophy and the relativist revolt against Western civilization Dressage masterclass Cobol for Students, Fourth Edition The Twelve Prophets (Vol. 2): Micah, Nahum, Habakkuk, Zephaniah, Haggai, Zechariah, Malachi (Berit Olam s Loomings: dread in west Texas Ear, nose, throat, and dental disorders The Song of the Cardinal and at the Foot of the Rainbow Treatise and hand-book of orange culture in Florida. Part 3 : What Aneth saw. Mad hatters holiday (1973 by Peter Lovesey States of extremes Les mills pump fitness guide New Yorks 50 best wonderful things to do at the holidays Better place to live