CHALMERS UNIVERSITY OF TECHNOLOGY | UNIVERSITY OF GOTHENBURG

Exam

DAT277/DIT053: Energy-Aware Computing
DAT278/DIT054: Sustainable Computing

Wednesday 08 April 2020, 08:30-12:30 (+10min tolerance)

Examiner and contact person during exam:
Pedro Petersen Moura Trancoso (contact through canvas email)

Grading intervals:
The maximum grade for this is exam is 100 points
DAT277/278 (Chalmers): U < 40p ≤ Grade 3 < 60p ≤ Grade 4 < 80p ≤ Grade 5
DIT053/054 (GU): U < 40p ≤ Grade G < 75p ≤ Grade VG

Examination solution and review:
Exam solution will be posted in canvas within 24h after the exam has finished. A canvas announcement regarding the exam review.

Instructions:

- **ANSWERS**: All answers need to be typed in the space provided. Different questions need to start in a new page. Additional material may be added to your typed answers as pictures inserted in the same space with the typed text. (pictures can be used only for diagrams not for text)
- **LANGUAGE**: All answers should be written in English.
- **SUBMISSION**: Submit your answers in a SINGLE PDF file. Name your file "<course code>-<first name>-<last name>.pdf" (e.g. for John Smith taking DAT278 would be "DAT278-John-Smith.pdf"). You may submit many times during the exam time but only your LAST submission will be considered.
- **DEADLINE**: The deadline is 12:30, there is a 10 min tolerance but then the assignment closes at 12:40, no more submissions will be accepted after that. Submit before the deadline to avoid problems.
- **RESOURCES**: All available material is allowed. For any of your answers, if you use information from the internet, include the source (the web address).
- **PLAGIARISM**: Copying of material from any source is FORBIDDEN (you need to write the answers in your own words). Communication and exchange of information with other people during the exam period is FORBIDDEN. Exam answers will be tested for plagiarism.
- **COMMUNICATION WITH EXAMINER:** Before communicating with the examiner test the Frequently Asked Questions (FAQ) page which will contain the answers given to other students. This page will be continuously updated. If you still need to communicate with the examiner, do this through canvas email.
- **CONTINGENCY PLAN**: In case there is a failure in canvas during submission you will have to submit your file with the answers as an attachment to an email to the examiner ppedro@chalmers.se. Use the same email in case the canvas site is down during the exam period and you need to communicate with the examiner.

Question 1: Sustainability and Sustainable Computing (20 points)

(a) (10 points) From the point of view of "sustainability" compare the traditional hall exams that you have taken in the past compared with the home exam you are taking now. In your comparison discuss also (but not only) the computer resources used for each type of exam (consider the whole process, from creation – examination – grading). Which exam is more "sustainable"? Justify your answer.

For both exams the examiner needs to first prepare the exam on the personal computer. Internet resources may be used to get information for the exam questions as well as to store the exam file on the cloud.
After this, for the hall exam, the file sent by email to the person responsible for printing the exams. A certain number of exams are then printed on paper. At the exam date the exams are sent to the rooms and answer sheets are also provided for the students to write their answers. At the end exams are collected and sent back to the examiner for grading. The grading is done on the exam papers.
In contrast for the home exam, after the exam is prepared, an online assignment in canvas is prepared where the file is uploaded. The students download the file and fill in the answers on their computers and then upload the answers. After the submission deadline, the examiner has access to the submissions and can grade the answers on the system using the personal computer. The hosting system is located in a remote place in Sweden.
So, while the traditional exam heavily depends on paper the home exam depends heavily on using computers by each student for the 4 hours of the exam and computer by the examiner for grading, as well as cloud storage for all files. In order for the submissions not to be lost by accident, replicas are possibly used increasing the demands on the cloud resources. Even though it may seem that the traditional exam has more "waste", if recycled paper is used for the traditional exam as well as the exams are send for recycling after storage period then it is not necessarily that the home is more "sustainable" we need to also to be aware of how the energy that was used for all computers and cloud storage was produced (e.g. green energy from renewable sources).

Question 1: Sustainability and Sustainable Computing (20 points)

(b) (10 points) As you know, there are different types of computer systems – mobile, personal, server, data centers, etc. Given the current situation where most people are working from home, which types of computer systems are used, and which should be the focus for improving their efficiency in order to reduce the overall carbon footprint and why? Justify your answer.

For people working at home several systems are being used at the same time. People use their mobile phones for example to communicate with other friends or colleagues or check for the latest news or updates. Then people use personal computers to do their work locally and access resources on the Internet. Cloud and Data centers are used as service providers for many applications (e.g. Google docs) as well as storage for remote cloud storage (e.g. Box). There is no single system that is most important, efficiency should be improved at all levels and through all systems. On the one hand, the mobile phones may be the systems consuming the least power but are also the largest number of devices in the world so there is a multiplicative effect for the efficiency improved on each device. On the other hand, the data centers are the systems easier to control and improve since they are somehow centralized and thus easier to manage.

Question 2: Power, Energy and Efficiency (20 points)

(a) (10 points) Use the internet to find the power consumed (or power dissipation or TDP) by a CPU used in a low-end laptop (System LOW – power should be lower than 50W - e.g. Intel Celeron or AMD A) and a CPU used in a high-end server (System HIGH – power should be higher than 100W - e.g. AMD EPYC or Intel Xeon). Mention in your answer the CPU models, number of cores, frequency, L2 cache sizes, power consumption, and the web address for your sources. Then consider that we execute the same application on both systems and that on system LOW it takes 100s to complete its execution. Give a rough estimation of the execution time on system HIGH (estimate using the difference in frequency, cache, number of cores in the systems). Justify this estimate. Compare the efficiency metric EDP of system LOW and HIGH for the execution of this application. Which system is more efficient? Justify your answer.

System LOW: Intel Celeron J4015, 2 cores, 2GHz frequency, L2 cache 4MB, 10W power (https://en.wikipedia.org/wiki/List_of_Intel_Celeron_microprocessors)
System HIGH: AMD EPYC 7601, 32 cores, 2.2GHz frequency, L2 cache 512KB/core so 16MB, 180W power (https://en.wikipedia.org/wiki/Epyc)

Application A on system LOW takes 100s to execute. The frequency is almost the same, so it is just the number of cores the cache size. The cache for the Celeron is actually larger per core (2MB vs 512KB). If we assume that the application is compute bound and can scale with the number of cores we could estimate that for 16x number of cores we could have an improvement of roughly 10x and thus the execution on the HIGH system could be 10s.

Efficiency = EDP = Energy x Delay = Power x Execution x Execution

EDP(LOW) = 10 x 100 x 100 = 100000
EDP(HIGH) = 180 x 10 x 10 = 18000

The lower the EDP shows the most efficient system thus HIGH is much more efficient than LOW.

Question 2: Power, Energy and Efficiency (20 points)

(b) (10 points) There are different methods to measure power in a system. In the slides of DAT278 Lecture 4 it states that some of these methods are:

- Simulation,
- Power-plug meters,
- On-die thermal sensors,
- Performance-counter-based power and thermal estimates, and
- Thermal imaging

Choose two (2) of these methods, briefly describe them, give an example (give the web address of your source) and compare as to what are the advantages and disadvantages for each method when the goal is the evaluation of which existing processor would be best for a mobile phone. The evaluation would be done by running the same benchmark applications on the different processors. Justify your answer.

Simulator: McPAT is a processor power simulator (https://github.com/HewlettPackard/mcpat) which uses a configuration file with the characteristics of the processor and then McPAT has some mathematical models to determine the power consumption for the execution of different instructions.

Power-plug meters: Revogi Smart Meter Plug (https://www.revogi.com/smart-power/smart-meter-plug-eu-2/#section0) this is a device that you can place between the power supply and the device. It measures and reports on an app the instant power consumption and keeps a record of the power consumption over time.

None of these two is very appropriate for the task. The Simulator will only be usable if you have a good model for the existing processor. Then you can use the characteristics for the different processors and run the applications through McPAT in order to get the results you need for your comparison. The Simulator advantages is that it is a controllable environment and that you may change the configurations of the processor very easily to get different results. This is ideal for the development phase of the processors, not so much for existing processors.

The power-plug meter is not usable at all since the mobile phone executes on battery so it would be impossible to find out its consumption. Even when charging, the power consumption is not an indication of the use of the device but just of the charging of the battery. Thermal sensors based modeling would be a much better solution since we could execute real applications on the real hardware to perform the required evaluation.

Question 3: Technology and Circuits (20 points)

(a) (10 points) In the slides of DAT278 Lecture 5 we present examples of processors that operate under subthreshold voltages. Describe in your own words what the "subthreshold voltage" technique is. What are the advantages and disadvantages of this technique? Justify your answer. Find a research article on such a processor and provide the title and web address.

Subthreshold Voltage technique: This is a technique where the processor operates at a voltage that is below the regular threshold voltage. This requires the use of specially designed transistors using different materials and combining this with other techniques so that the transistor switching is still activated and detected properly.

The clear advantage is that when operating at such a low voltage, the power consumption becomes minimal. This is great for IoT devices with very limited battery or even devices that get their energy from harvesting (e.g. from heat).

The disadvantage is that at such low voltages, the static power increases dramatically. So, if regular materials and transistors were to be used then the static power disadvantage would strongly offset the dynamic power benefit.

A possible research paper: "Energy-Efficient Subthreshold Processor Design" (https://ieeexplore.ieee.org/document/4895693)

(one possible source: https://www.extremetech.com/computing/198285-new-microprocessor-claims-10x-energy-improvement-thanks-to-subthreshold-voltage-operation)

Question 3: Technology and Circuits (20 points)

(b) (10 points) Find in the internet and example of a multicore processor with 64 cores. One technique for improving the energy efficiency that can be used in multicore processors is the use of multiple voltage domains. Why do you think there is an advantage in having multiple voltage domains in a multicore processor? Give a practical example where this feature would be helpful. Do you believe that in the processor you have found there is one voltage domain per core, so 64 voltage domains? Justify your answer.

The AMD Ryzen Threadripper 3990X Processor (https://www.amd.com/en/products/cpu/amd-ryzen-threadripper-3990x) is a 64 core high-end processor.

When you have many cores on the same processor and different applications executing on that processor simultaneously, maybe each application has different requirements and thus it would be more efficient to have a different voltage for each core so that the power would be adapted at a fine grain to the requirements for each core. This would result in a much more efficient solution than operating all cores at the same voltage.

It is not clear from the description found for the processor above that there are multiple voltage domains or one voltage domain per core. It is highly improbable that that is the case since supporting multiple voltage domains is costly in terms of silicon space to implement isolation islands between the different domains. Also, it is difficult to maintain multiple voltage regulators. Thus, a more common solution is to group cores into clusters and then provide one voltage domain per cluster. For example, the 64 cores could be divided into 4 clusters of 16 cores each and then have a voltage domain per cluster. So demanding applications could all execute on the high voltage cluster and non-demanding applications could run on the low voltage cluster.

Question 4: Dynamic Power (20 points)

(a) (10 points) DVFS can be implemented at different levels:

- Hardware level
- System level
- Application/Programmer level

Explain in your own words what DVFS is. Discuss the advantages and disadvantages of the implementation at each level. Which level is best and why? Justify your answer.

DVFS is a technique where voltage and frequency are changed as to adapt the power consumption of the processor to the requirements of the application(s) running on that same processor. Different voltage-frequency states are usually available for the processor operation.

Hardware-level DVFS can be implemented in the low-level hardware devices so to control the voltage and frequency at the transistor level. This is very efficient in the sense that all components could potentially be controlled for optimal overall execution. This would be difficult to control and coordinate with some many hardware components. Also, the hardware does not have a global view of the requirements, so it is difficult to estimate the future need, just possible to react based on past knowledge.

System-level DVFS is done usually at the Operating System (OS) level. The OS has a much better system view of the requirements and thus can better orchestrate changes in voltage and frequency, as well as being able to combine the scheduling of the execution with the different voltage-frequency states or proactively change voltage-frequency state. It still does not have a perfect view of the application demands and the voltage-frequency changes are much slower than at the hardware level.

Application-level DVFS is done at the application by the programmer. As such, the programmer has a perfect knowledge of the application requirements and can effectively request the best matching voltage-frequency state for a certain task. In this case, while the decisions will be the best for that particular application, since this ignores the co-execution with other applications on the same system, those decisions may not be optimal at the end. In addition, voltage-frequency state changes at this level are the costliest since the request needs to traverse all the system from up to down. Consequently, changes cannot be done very frequently.

As it is clear there is no best solution. Each solution has its pros and cons. It depends on the application and the system what works best but I would say that the best would be a combination of all techniques in a collaborative way to solve the problem and achieve the best overall efficiency.

Question 4: Dynamic Power (20 points)

(b) (10 points) In the AMD Ryzen 7 1700X processors, the L3 is a 16 MB, 16-way set-associative cache. Considering energy-efficiency, why is it beneficial to have 16 ways in such a cache? What type of misses are reduced with the large associativity? What is the problem though? This problem is addressed by the "way-prediction" technique. Explain how way prediction works. What are its advantages and disadvantages? Justify your answer.

The benefit of a large degree of associativity is that a certain data element has more opportunities to be placed in the cache, thus reducing the conflict misses. The problem is that there is an increased metadata storage since for every set, there are as many tags as the number of ways. This means that the cache needs to have more space for non-data, thus is more costly in terms of space. Also, all ways need to be activated for the tags to be compared and thus there is an increase in energy consumption. It is important to note that at best the data will be found in one of the 16 ways, so in the best case there is a 15/16th waste of tag search.
Way prediction is a technique that can be used to just test a single tag out of the 16 so that energy is saved. So, the regular cache access is combined with a good predictor that can predict in which way the data can be found. If the predictor is predicting correctly, then the energy saved is considerable. If there is a miss prediction, then all tags need to be checked and thus there is a penalty in terms of extra energy for the second access as well as latency.

Question 5: Heterogenenity and Data Centers (20 points)

(a) (10 points) Find a definition for “Dark Silicon” and include it in your answer with the reference (web address) to where you go it. Describe in your own words what it is. Which problem is being created by this effect? How is the utilization of a processor affected by Dark Silicon? Describe one solution that the researchers have proposed to overcome this issue (the title of Q5 may give you a hint!). Give an example of an existing product (e.g. processor) and justify your answer.

“In the electronics industry, dark silicon is the amount of circuitry of an integrated circuit that cannot be powered-on at the nominal operating voltage for a given thermal design power (TDP) constraint. This is a challenge in the era of nanometer semiconductor nodes, where transistor scaling and voltage scaling are no longer in line with each other, resulting in the failure of Dennard scaling. This discontinuation of Dennard scaling has led to sharp increases in power densities that hamper powering-on all the transistors simultaneously at the nominal voltage, while keeping the chip temperature in the safe operating range. According to recent studies, researchers from different groups have projected that, at 8 nm technology nodes, the amount of Dark Silicon may reach up to 50–80% depending upon the processor architecture, cooling technology, and application workloads. Dark Silicon may be unavoidable even in server workloads with abundance of inherent client request-level parallelism.” (in https://en.wikipedia.org/wiki/Dark_silicon)

So dark silicon describes the situation where we have many transistors in a chip but that given power constraints, we are not able to switch (i.e. use) all transistors simultaneously. The power density is the cause for this effect, many components are packed very densely and there is not enough space to dissipate the power produce by their switching. Consequently, we can use only parts of the whole chip in order to avoid problems with high power dissipation thus reducing the maximum possible utilization.

Researchers have come up with different solutions. One is to reduce the voltage considerably so that the power dissipated is reduced considerably and then all transistors can again be used simultaneously. This would though lead to sub-threshold solutions.

Another solution is to make processors heterogeneous so that there is different hardware specific for each different application and when a certain application executes, its tasks are directed to that specific efficient hardware while the rest can stand idle. On example is the Apple A13 Bionic (https://www.macworld.com/article/3442716/inside-apples-a13-bionic-system-on-chip.html) which has different types of integrated processing engines, some for graphics (GPUs), some for Artificial Intelligence (AI), some for security, etc. This is a combination of many different very specific but efficient components on the same package thus giving the overall efficiency.

Question 5: Heterogenenity and Data Centers (20 points)

(b) (10 points) Microsoft has decided to use FPGAs in their Azure Data Center. How are the FPGAs used in Microsoft's Project Brainwave? Which would be other alternative hardware components that are used for the same task? What is the benefit of using FPGAs as compared to those alternatives? Justify your answer.

In Microsoft's project Brainwave (https://www.microsoft.com/en-us/research/project/project-brainwave/), FPGAs are used to implement Neural Processing Units (NPUs) to efficiently solve, for example, image classification problems using Deep Neural Networks (DNNs). Instead of FPGAs we could use GPUs or TPUs. The latter are dedicated hardware for Machine Learning from Google (https://cloud.google.com/tpu).

The advantage of using FPGAs is that the efficiency is very high for FPGA implementations in most cases. In addition, the fact that FPGAs can be re-configured allows them to adapt to changes. For example, you may upload to an FPGA a particular model for ML inference and then after re-training the model you may upload a new more optimized version just by reconfiguring the hardware. This gives a large flexibility to the system.