DOWNLOAD PDF A TRAITOR TO MEMORY, (2 PART SET: PART 1 12 TAPES PART 2 9 TAPES)

Chapter 1 : MiniDV Camcorder Tapes for sale | eBay

UK VHS release (two-episode tapes, CIC Video): Volume 79, 9 May In feature-length form, as part of the UK VHS release Star Trek: The Next Generation - The Full Length TV Movies: Volume 9, catalogue number VHR , 8 May

Thank God for that. On the planet Helotrix, a Federation Colonel named Quute is interrogating an insurrectionist named Igin. Why are the Federation obsessed with eye patches? Oh, come on, Sue, you must remember him. Did he give birth to the blonde one from Peep Show? Because he looks just like him. I keep expecting them to break into song. At least the location gets the thumbs-up. All the planets were beginning to look the same to me, but this one is weird enough to be truly alien. I really like it. I mean, has anybody come looking for them recently? Avon could just change his name and move somewhere quiet if he really wanted to. What trouble could he possibly cause? And why would he bother? Why would any of them bother? Back on Helotrix, Colonel Quute plays a game of chess as he casually wipes Igin off the face of the planet. That rebel was just a pawn in his game, right? Can you see what they did there, Neil? So they eat, drink, play chess and kill people in one single room. Does he sleep in there when his shift is over? Has he got a sofa bed tucked away somewhere? At least the Liberator had more than one set. I did say all this two weeks ago. And is Soolin just paid to lounge around looking pretty? Avon has completely lost the plot. I expected a lot more than this from Robert Holmes. Blah blah Magnowhatsit Terminal blah blah blah. Leitz, General Federer and President-Elect Practor discuss the problem of the rebels and how best to deal with them. Yak yak bloody yak. With Robert Holmes you usually get some decent lines but this is tedious, Neil. Hunda swims into the city under the cover of night. Dayna and Tarrant teleport to the surface of Helotrix and immediately lose contact with the ship. The name Slave makes me feel uncomfortable. They should change it. What would you suggest? The kids would have loved that. This is supposed to be an adult drama, Sue. Which reminds me, Tarrant let me stroke him the other day. The President-Elect is shown to his quarters, where a painting of the Supreme Empress hangs on one of its wall. The only surprising thing about this is that it took this long for Servalan to show up again. Of course it is, Neil. What were all those people doing out there in the first place? Were they having a fag break? That would explain the smoke, I suppose. And why is Orac telling them what to do? Is Orac the traitor? I like their guns. Everyone should have one. Twenty minutes laterâ€¦ Sue: Oh look, a soldier just fell into a pond. How many is that now? And do you know what, Neil? The Helots are no match for the rebels. The Helots are swiftly dealt with. I bet he sat by his phone waiting for his agent to call when Tom Baker left. He would have been perfect. Better than Peter Davison? Imagine thatâ€¦ A Federation guard pacifices a Helot. The poor sod is left with a creepy grin on his face. Avon refuses to leave Tarrant and Dayna to their fate, which prompts a visibly annoyed Vila to claim that Blake would have been proud of him. Avon is just as crazy as Blake used to be, if not more so. At least Blake had a reason to fight the Federation, after everything they did to him. He should be out robbing banks. A mysterious figure kills the President-Elect in cold blood. Why are they trying to hide it? Tarrant disguises himself as a Helot. Are you saying you fancy Steven Pacey as well? It appears that Leitz has been helping the rebels all along. Is it over yet? Leitz suggests that the rebels storm the city via an abandoned monorail system. This what happens when you sub-contract. The real Orac would have covered his tracks better than that. Leitz confers with a mysterious, gravely voice over an intercom system. How much more of this is there? Tarrant and Dayna encounter a mad professor in a wheelchair, which certainly livens things up a bit. Forbus plans to blow himself and Commissioner Sleer to bits, the first chance he gets. Oh, and there are some drugs in it as well. Avon decides to take the Scorpio below cloud level to avoid detection. Why does everything have to end in a row? Commissioner Sleer shows her face. Have you guessed who it is yet? Tarrant and Hunda fight the Federation. This is pretty bad. The fight, the wobbly scenery, the music, the fake rocks, the fact that this on video instead of film, the music. Did I mention the music? Sue is also adamant that this episode drops the C-word before the watershed. Servalan quickly dispenses with Forbus. She must have been one of the most famous people who ever lived. Servalan and Leitz kiss, but before Leitz can catch his breath, Servalan stabs him in the back. What is it with this show and violent snogging? It happens all the time. I needâ€¦ to kill her myself. Well that was shit. Worse than that, it was boring shit. What else do you want me to say? I used to know the Production Manager â€” Rosie Crowson. She taught me everything I know about TV production. The only good thing about that episode was Servalan. At least it gives Avon something to do.

DOWNLOAD PDF A TRAITOR TO MEMORY, (2 PART SET: PART 1 12 TAPES PART 2 9 TAPES)

Chapter 2 : McCain and the POW Cover-up [Part 2] (July 20,)

Eddington also states that Cal Hudson, played by Bernie Casey in "The Maquis, Part I" and "The Maquis, Part II", has been killed by the Cardassians, thus providing a degree of closure for Sisko.

This probability distribution is selected to improve traceability for a particular size of a coalition of attackers. At least one symbol for each file-segment variation is then distributed based on the selected probability distribution. The transition of many types of media from analog to digital offers new advantages to the consumer in quality and flexibility. Also, there is an increasing use of global distribution systems such as the Internet for the distribution of digital assets, including music, film, computer programs, photographs, games and other content. These trends have made it easy to produce and distribute flawless copies of content by content providers. Unfortunately, there is also a concurrent increase in the unauthorized copying, or pirating, of digital content, which has caused considerable economic losses to content providers. Effective countermeasures are important to the viability of businesses engaged in the distribution of digital media. Piracy is a major concern and expense for content providers. CPRM is a technology developed and licensed by the 4C group, comprising IBM, Intel, Matsushita, and Toshiba, to allow consumers to make authorized copies of commercial entertainment content where the copyright holder for such content has decided to protect it from unauthorized copying. In the AACS content protection system, devices such as DVD players are assigned a set of keys and a common key is used to encrypt the content. A pirate attack in this system may occur when the attackers redistribute the common content encrypting key or the plain content to avoid being identified. This type of an attack is called an anonymous attack. In an anonymous attack, an attacker, or group of attackers, tries to hide their secret device keys and operate anonymously. In this attack, the attackers instrument their devices and collude to build a pirate copy of the decrypted plaintext content or the decryption key itself. The attackers can then redistribute either the plaintext content, or the decryption key. To do traitor tracing for anonymous attacks, content may be divided into multiple segments and some of the segments may be chosen to have multiple variations. A digital watermark is one way to build these variations. More importantly, those variations are not only differently watermarked, but also differently encrypted. During playback, each device can only decrypt exactly one variation at each segment. The differently watermarked and encrypted variations effectively build different content versions. Each different playback path becomes one version. The recovered pirated variation encrypting keys, or the movie version, can be linked back to the actual devices i. There are some practical issues with the above-described traitor-tracing system. First of all, because the variations take extra space on the disc bandwidth during communication, the number of variations cannot be large. However, in practice, the number of devices a system needs to accommodate may be very large, e. These are conflicting requirements. The inner code assigns the variation for each segment inside the content, which may be a movie. This assignment effectively creates multiple movie versions, each version becoming a symbol for the outer code assignment. The outer code assigns the movie versions symbols among a sequence of movies. This assignment solves the extra-bandwidth requirement by having a small number of variations at each segment, while still managing to support a large number of devices. A second practical issue relates to the actual traitor detection. The problem is that attackers collude in the attack and may mislead the tracing agency to erroneously incriminate innocent devices. The collusion attack creates an inherent difficulty in terms of tracing. After the above-described practical assignment is done, a straightforward approach to detect colluders might be to score every device and incriminate the highest scoring devices. In some prior systems, more efficient tracing algorithms are employed which use a set-cover algorithm to detect coalitions of pirates all together instead of one by one. In these systems, when the number of traitors becomes large, the traceability decreases. Hence, prior-art systems using a set-cover tracing algorithm may work fine when the number of traitors is smaller than q, with q being the number of symbol variations. When the number of traitor exceeds q, the traceability degrades significantly. When the number of traitor reaches q log q, where the coalition gets to know every symbol, the scheme may be nearly broken. According to another embodiment of the present invention, a method comprises: According to a further embodiment of the present invention, a system comprises: According to another embodiment of the present invention, a computer-program product for traitor tracing, comprises: The present invention can perform traitor tracing where the number of traitors is relatively large, for example larger than q, the number of symbol variations. In this way, the present invention raises the tracing limit on the number of traitors as compared to prior traitor-tracing systems. The present invention, in some embodiments, may also adaptively assign the symbols in the inner code so as to maximize the traceability for a traitor size. Traditionally, for a segment in the movie to have multiple variations, the symbols assigned to the variations are equally likely to appear. With the present invention the equal probable assignment of the symbols is modified so that any kind of probability assignment, including nonuniform probability assignments, may be used. The present invention may achieve significant advantages over the prior art. First, it enables the tracing agency to adaptively assign the symbols with a particular probability, as needed, in order to achieve superior even potentially maximum traceability for any traitor size. Secondly, with the unequal probability assignment of the symbols, the traitor size limit is pushed much further. For example, the present invention can achieve good traceability when traitor size exceeds q, and it continues doing well even after the coalition size reaches q log q. To do traitor tracing for anonymous attacks, content is divided into multiple segments and some of the segments are chosen to have multiple variations. The inner code assigns the variation for each segment inside a piece of content, e. This assignment effectively creates multiple movie versions, and each version becomes a symbol for the outer code assignment. The outer code assigns the movie versions symbols to players for a sequence of movies. In more detail, referring now to FIGS. System 10 comprises a software programming code or a computer program product that is typically embedded within, or installed on a media 15 and a media player Alternatively, system 10 can be saved on a suitable memory or storage medium such as a diskette, a CD, a DVD, a hard drive, or like devices. The media player 20 can access a web service provider 25 through a network The media player 20 comprises software that allows the media player 20 to interface securely with the web service provider The media player 20 is connected to network 30 via a communications link 35 such as telephone, cable DSL, satellite link, etc. The web service provider 25 is connected to the Internet through a communications link The media player 20 downloads content from the web service provider 25 and records the content on the media Alternatively, media 15 may be prerecorded with content such as, for example, movies, audio files, video games, computer programs, or any other type of electronic file. Media 15 may be played on the media player Content on media 15 comprises one or more files. Files may comprise any kind of digital data sequence, including but not limited to text, audio, images, video, music, movies, multimedia presentations, operating systems, software applications, and cryptographic keys. In broad terms, file includes a beginning , an end , and a span of data. Files may be of any size and may be distributed by any means, including but not limited to computer networks, satellite networks, cable networks, television transmissions, and various physical storage media e. Files may be broadcast in groups in a substantially continuous sequence. The present invention may be applied to any digital content subject to one-to-many distribution. Movie rental boxes are one application. Another example may be for operators of a web server generally referred to as a digital rights manager that sells copyrighted content such as music, or other material stored in a subscription database, who may not want to encrypt or otherwise process files on the fly because of the computational expense involved. Similarly, such a server can neither individually tailor nor store a complete copy of every file it transmits. Referring now to FIG. For clarity, only three critical file segments are shown; a typical number may be approximately Not all data in a file need to be protected to the maximum-possible level of security; bandwidth can be conserved by selectively applying different levels of security to the most valuable portions of a file For example, when the file comprises a movie, each movie may have scenes that are each absolutely essential for the movie to be acceptable to any audience. All critical file segments in a file must therefore be properly processed for the file to be commercially desirable. Embodiments of the present invention may select five-second scenes in a typical movie as critical file segments, but critical file segments of varying length are also encompassed by the present invention. The critical file segments are not necessarily equally distributed throughout a given file, in fact the critical file segments are preferably especially selected based on the contents of the file, possibly by human editors. In the case of executable software files, automated tools may identify critical file segments according to a measured execution frequency. Referring now to FIGS. For clarity, only four file segment variations are shown for each critical file segment; a typical number may be approximately Each file-segment variation is simply a copy of the particular corresponding critical file segment that has been differently watermarked and differently encrypted. Each entire file is also typically watermarked and encrypted in a broadcast encryption system. Each file-segment variation is identified by a text designation in this application e. The number of critical file segments and the number of file segment variations employed may depend on the properties of the file and its audience. For movies, one could select a single critical file segment and have several hundred file-segment variations; however, attackers may simply choose to omit that single critical file segment in a pirated copy of the file, in hopes that viewers would not find such a glitch to be overly annoying. A pirated movie with, say 15 missing critical 5-second scenes is probably going to be too annoying to any viewer for it to be of any commercial value. Thus, the illegally-broadcast movies are either substantially disrupted or the attackers must incorporate some of their file segment variations, which would facilitate traitor tracing. While the number of critical file segments and the number of file-segment variations may be kept constant for each file, modifying either number according to an estimated piracy likelihood for a given file is also within the scope of the invention. The number of file segments and the number of file-segment variations will determine the amount of bandwidth overhead or, alternately, the increased size of the broadcast version of the file. The augmented file is the version of the original file that will actually be broadcast. Each intended receiver of the broadcast of a group of files requires augmentation-selection information to choose a particular combination of file-segment variations for each particular file. For example, a movie player must know, for each movie, which set of variations to plug into the spaces where critical scenes existed in the original movie. The particular arrangement of unmodified file content and file segment variations within the augmented file shown is not critical but merely intuitive. The techniques employed by the present invention facilitate traitor tracing in a commercially-viable i. If different boxes are assigned different combinations of file-segment variations to use, an analysis of a pirated file can help determine which boxes were used as part of an anonymous attack.

DOWNLOAD PDF A TRAITOR TO MEMORY, (2 PART SET: PART 1 12 TAPES PART 2 9 TAPES)

Chapter 3 : â€˜The Handmaidâ€™s Taleâ€™ Recap: Season 2, Episode 3 â€” â€˜Baggageâ€™ | TVLine

This feature is not available right now. Please try again later.

These IBM tabulating machines from the s used mechanical counters to store information A portion of a core memory with a modern flash SD card on top 1 Megabit chip â€” one of the last models developed by VEB Carl Zeiss Jena in Early computers used relays , mechanical counters [3] or delay lines for main memory functions. Ultrasonic delay lines could only reproduce data in the order it was written. Drum memory could be expanded at relatively low cost but efficient retrieval of memory items required knowledge of the physical layout of the drum to optimize speed. Latches built out of vacuum tube triodes , and later, out of discrete transistors , were used for smaller and faster memories such as registers. Such registers were relatively large and too costly to use for large amounts of data; generally only a few dozen or few hundred bits of such memory could be provided. The first practical form of random-access memory was the Williams tube starting in It stored data as electrically charged spots on the face of a cathode ray tube. Since the electron beam of the CRT could read and write the spots on the tube in any order, memory was random access. The capacity of the Williams tube was a few hundred to around a thousand bits, but it was much smaller, faster, and more power-efficient than using individual vacuum tube latches. Developed at the University of Manchester in England, the Williams tube provided the medium on which the first electronically stored program was implemented in the Manchester Baby computer, which first successfully ran a program on 21 June It became a widespread form of random-access memory, relying on an array of magnetized rings. Since every ring had a combination of address wires to select and read or write it, access to any memory location in any sequence was possible. Magnetic core memory was the standard form of memory system until displaced by solid-state memory in integrated circuits, starting in the early s. Dynamic random-access memory DRAM allowed replacement of a 4 or 6-transistor latch circuit by a single transistor for each memory bit, greatly increasing memory density at the cost of volatility. Data was stored in the tiny capacitance of each transistor, and had to be periodically refreshed every few milliseconds before the charge could leak away. Prior to the development of integrated read-only memory ROM circuits, permanent or read-only random-access memory was often constructed using diode matrices driven by address decoders , or specially wound core rope memory planes. In SRAM, a bit of data is stored using the state of a six transistor memory cell. As this form of memory is less expensive to produce than static RAM, it is the predominant form of computer memory used in modern computers. Both static and dynamic RAM are considered volatile, as their state is lost or reset when power is removed from the system. By contrast, read-only memory ROM stores data by permanently enabling or disabling selected transistors, such that the memory cannot be altered. These persistent forms of semiconductor ROM include USB flash drives, memory cards for cameras and portable devices, and solid-state drives. Memory cell Main article: Memory cell computing The memory cell is the fundamental building block of computer memory. The memory cell is an electronic circuit that stores one bit of binary information and it must be set to store a logic 1 high voltage level and reset to store a logic 0 low voltage level. The value in the memory cell can be accessed by reading it. This means that SRAM requires very low power when not being accessed, but it is expensive and has low storage density. A second type, DRAM, is based around a capacitor. Charging and discharging this capacitor can store a "1" or a "0" in the cell. However, the charge in this capacitor slowly leaks away, and must be refreshed periodically. Within the RAM device, multiplexing and demultiplexing circuitry is used to select memory cells. Typically, a RAM device has a set of address lines A An, and for each combination of bits that may be applied to these lines, a set of memory cells are activated. Due to this addressing, RAM devices virtually always have a memory capacity that is a power of two. Usually several memory cells share the same address. Often the width of the memory and that of the microprocessor are different, for a 32 bit microprocessor, eight 4 bit RAM chips would be needed. Often more addresses are needed than can be provided by a device. In that case, external multiplexors to the device are used to activate the correct device that is being accessed. Memory hierarchy Main article: Memory hierarchy One can read and over-write data in RAM. Many computer systems have a memory hierarchy consisting of processor registers , on-die SRAM caches, external caches , DRAM , paging systems and virtual memory or swap space on a hard drive. This entire pool of memory may be referred to as "RAM" by many developers, even though the various subsystems can have very different access times , violating the original concept behind the random access term in RAM. Even within a hierarchy level such as DRAM, the specific row, column, bank, rank , channel, or interleave organization of the components make the access time variable, although not to the extent that access time to rotating storage media or a tape is variable. The overall goal of using a memory hierarchy is to obtain the highest possible average access performance while minimizing the total cost of the entire memory system generally, the memory hierarchy follows the access time with the fast CPU registers at the top and the slow hard drive at the bottom. In many modern personal computers, the RAM comes in an easily upgraded form of modules called memory modules or DRAM modules about the size of a few sticks of chewing gum. These can quickly be replaced should they become damaged or when changing needs demand more storage capacity. In addition to serving as temporary storage and working space for the operating system and applications, RAM is used in numerous other ways. Virtual memory Main article: Virtual memory Most modern operating systems employ a method of extending RAM capacity, known as "virtual memory". When the system runs low on physical memory, it can " swap " portions of RAM to the paging file to make room for new data, as well as to read previously swapped information back into RAM. Excessive use of this mechanism results in thrashing and generally hampers overall system performance, mainly because hard drives are far slower than RAM. RAM disk Main article: A RAM disk loses the stored data when the computer is shut down, unless memory is arranged to have a standby battery source. The ROM chip is then disabled while the initialized memory locations are switched in on the same block of addresses often write-protected. This process, sometimes called shadowing, is fairly common in both computers and embedded systems. Depending on the system, this may not result in increased performance, and may cause incompatibilities. For example, some hardware may be inaccessible to the operating system if shadow RAM is used. On some systems the benefit may be hypothetical because the BIOS is not used after booting in favor of direct hardware access. Free memory is reduced by the size of the shadowed ROMs. The technologies used include carbon nanotubes and approaches utilizing Tunnel magnetoresistance. There are two 2nd generation techniques currently in development: Whether some of these technologies can eventually take significant market share from either DRAM, SRAM, or flash-memory technology, however, remains to be seen. Since , " solid-state drives " based on flash memory with capacities exceeding gigabytes and performance far exceeding traditional disks have become available. This development has started to blur the definition between traditional random-access memory and "disks", dramatically reducing the difference in performance. Some kinds of random-access memory, such as "EcoRAM", are specifically designed for server farms , where low power consumption is more important than speed. An important reason for this disparity is the limited communication bandwidth beyond chip boundaries, which is also referred to as bandwidth wall. Given these trends, it was expected that memory latency would become an overwhelming bottleneck in computer performance. Intel summarized these causes in a document. Secondly, the advantages of higher clock speeds are in part negated by memory latency, since memory access times have not been able to keep pace with increasing clock frequencies. Third, for certain applications, traditional serial architectures are becoming less efficient as processors get faster due to the so-called Von Neumann bottleneck , further undercutting any gains that frequency increases might otherwise buy. The End of the Road for Conventional Microarchitectures" [15] which projected a maximum of A different concept is the processor-memory performance gap, which can be addressed by 3D integrated circuits that reduce the distance between the logic and memory aspects that are further apart in a 2D chip. Multiple levels of caching have been developed to deal with the widening gap, and the performance of high-speed modern computers relies on evolving caching techniques.

DOWNLOAD PDF A TRAITOR TO MEMORY, (2 PART SET: PART 1 12 TAPES PART 2 9 TAPES)

Chapter 4 : Blaze of Glory (episode) | Memory Alpha | FANDOM powered by Wikia

Cardboard cassette tape with a portable USB memory stick included. Contains a 1GB memory capacity that holds minutes of music to create the perfect playlist for your loved ones. Cassette tape comes with blank section to inscribe the title and artist from the tracks on your special playlist.

Jake also likes it until Benjamin reveals that the sauce is puree of tube grubs , saying its only fair they try Ferengi food. Sisko turns to Nog to ask about his job with security. Jake whispers something and it is revealed that Nog has problems with the Klingons who ignore him because he is a Ferengi. Then, General Martok enters the room, interrupting the dinner and demanding to talk to Sisko alone. He explains to Sisko that his ship has intercepted an encrypted message from the Maquis. This is a surprise to both men, who believed they were completely wiped out by the Dominion. The message states that missiles have been launched heading for Cardassia. Sisko knows that if Cardassian citizens are killed by Human terrorists, the Dominion will launch a counter strike that will surely start a war. Worf states they need more information to narrow the search, and Sisko says he has an idea. Captain Sisko visits ex-commander Michael Eddington in a Federation prison. Eddington denies that the message about the missiles involves him. He says he has nothing to do with the Maquis anymore, as they are all dead, and they cannot be reborn. Eddington insists that the missiles are completely undetectable, and that nothing can be done to stop them. Sisko offers Eddington a path to eventual freedom if he will help find and destroy the missiles, but Eddington insists that he will be perfectly happy to wait in his cell until the missiles hit their target and the Dominion destroys his prison in their inevitable counter-strike. Back on the station, Quark is in the infirmary treated by Dr. Julian Bashir and questioned by Kira Nerys and Odo. It appears that Morn attacked Quark with a barstool. Quark claims it was unprovoked, but Kira and Odo reveal that Quark told Morn that Deep Space 9 is the first target if the Dominion attacks â€" and if that happens, there is no chance of survival. Kira wants to send the Defiant, but Sisko tells her he is not alone, and he reveals Eddington on board the runabout. Eddington is clearly affected by this, and sits down with his back to Sisko in a contemplative mood. Two Klingons are trading head blows for fun, with Martok watching and calling for more wine. Nog waits for the Klingons to get louder so he can send them to the brig for disturbance. When they exceed seventy decibels , Nog laughs and prepares to arrest them, but he falls from his chair and the Klingons laugh at him. He takes off his handcuffs and leaves piloting the ship while he gets a raktajino. This time, Eddington plans a Maquis trick to realign the impulse flow regulators while the engine is active. Sisko is alarmed, since this is a dangerous procedure, but Eddington says the plasma exhaust can be ignited and either destroy the enemy ships or at least mask their trail. Sisko agrees and gets in the crawlspace. Sisko, now with a head wound, comes out and berates Eddington for not keeping the ship steady, then ordering him to proceed to the missile base. Nog has a conversation with Martok On the station, Jake insists to Nog that, if he has a girl over, he must call before entering their quarters. Martok bellows that the young Ferengi is either courageous or a fool to threaten him with arrest. Nog trembles, but stands his ground. Eddington pilots the runabout to Athos IV.

DOWNLOAD PDF A TRAITOR TO MEMORY, (2 PART SET: PART 1 12 TAPES PART 2 9 TAPES)

Chapter 5 : USB2 - Adaptive traitor tracing - Google Patents

Hi everyone! I'm Zajcu37, 21 years old boy from Poland. My channel is devoted to the animation of the Fnaf and other games (in the future). Based on the Five.

Need to catch up? But then June thinks a minute and realizes that someone had noticed a disturbing cultural shift: Her reverie is interrupted by Nick, who has woken up and is preparing to leave. He tells her that she might be on her way soon. She refuses to move along the underground handmaid railroad without her daughter, but Nick counters that her getting out safe is better for everyone. He drops her at a warehouse where old street signs are stored side note: The man definitely thinks this is a bad idea, but he brings June to his home. In another flashback, June remembers watching the indoctrination slide show at the Red Center, right after she and Moira were made handmaiden. In bed that night, the pair wonders how Gilead got her, given that all the clinics destroyed their records. So she rides to a stop, then gets outâ€¦ and jumps into the woods the moment the Guardians policing the pathway near the train look away. Eventually, after crossing many fields in what are probably terribly uncomfortable shoes, June makes it to the airfield and waits until sundown for the arrival of her plane. When it does, the pilot is wary of her until she shows him her torn ear; he lets her scurry into the cargo hold, along with a former driver of a Commander. The pilot is dragged out and shot in the head. The driver, who took a bullet even before the plane came to a stop, is pulled out, too. They eventually grab her. He was in the U. Sometime later, Moira is somberly hanging out at a club when a woman catches her eye. They wind up in the bathroom, where Moira â€” how shall I put this? But when the woman offers to reciprocate, Moira shuts her down. Right before she leaves, they make awkward introductions. What did you think of the episode? Sound off in the comments!

DOWNLOAD PDF A TRAITOR TO MEMORY, (2 PART SET: PART 1 12 TAPES PART 2 9 TAPES)

Chapter 6 : Random-access memory - Wikipedia

By dividing 25, into 1 million (1,,), I get a quotient of 40, so a pile of new $ bills 40 inches tall would equal$1,,, a stack about up to the beltbuckle of a six foot tall man.

Edit The USS Enterprise-D is hit on the port nacelle with a barrage of fire from the mercenary vessel where Riker is being held prisoner. The same is true when the starboard nacelle is hit. Troi believes Riker has done something to their weapons, which Data , as acting captain, believes it to be true and they must play along. The Enterprise simulates serious damage and fires back with only minimal phaser power. This, along with the antimatter containment units buckling forces Baran to order the mercenary ship to withdraw. Data orders the crew to let the ship go, assuming that this is what Riker wants. Act One Edit Data is intent on holding still and investigating, rather than send ships to the other two likely planets to be attacked. On the mercenary ship, Riker and Picard stage a heated argument. Baran says he might like to continue to work with Riker, and tells Riker that he wants him kill Picard. Act Two Edit On the Enterprise, Data determines that the message is a flight plan, indicating the mercenary ship is headed to the Hyralan sector. Worf believes that it would take them at least 15 hours to be there since their maximum warp factor is 8. Data decides to beat them there, which Worf responds in an irritated tone, which catches both La Forge and Data off guard. Data brings Worf to the ready room and berates him over his performance as first officer , most notably: Worf states that he has always felt free to voice his opinions that would differ from Picard or Riker. Data accepts that but in those instances, he was acting as security chief. Data explains the role of the second-in-command is to carry out the orders of the commanding officer. In this case, the latter is Data. Data states that if Worf is not capable of carrying out this role, he will return him to tactical and assign La Forge as first officer instead. Worf elects to remain as first officer, and Data accepts it on the condition he conforms to the guidelines. He also apologizes if this berating has ended their friendship. Worf takes responsibility for it and asks if this incident be overlooked to continue their friendship. Data accepts it and Worf leaves. Once alone, Data simply straightens his uniform. Riker reveals that they are headed to the Hyralan sector to pick up the artifact, which will be Romulan in origin, from a Klingon transport ship. Picard, however, states that the artifacts are actually all Vulcan in origin. Picard discusses mutiny with Narik and is repulsed since he assumes Picard would want to be captain. He is favorable to Tallera, though. Then, Tallera bursts in and pulls a phaser on Picard, demanding to be told his true identity. The artifacts being sought are fragments of the mythical Stone of Gol , a powerful weapon she believes is being sought by a Vulcan isolationist group. She explains that the stone is a psionic resonator , a device that focuses and amplifies telepathic energy. If the stone were reassembled and used by a trained telepath, the weapon could eliminate the entire Vulcan council with a single thought. Meanwhile, in the Hyralan sector, the Enterprise is surprised to encounter a Klingon ship, Toron-class. Suspicious, they seek options on how to bring the craft over since the treaty with the Klingons gives its pilot, Koral free movement in Federation space. Worf suggests the treaty does stipulate taking the ship in for health and safety reasons. Koral is greeted by Dr. Crusher and Worf to conduct the "inspection". Koral scoffs at the notion but Beverly carries out the inspection, nervously trying to avoid any attempt at their real objective. When the mercenaries learn what has happened, Baran reluctantly prepares a team to board the Enterprise. Picard balks at the notion and asks how many security officers would serve on such as ship as that. Dismissing him, Baran orders Riker to lead the team to invade the Enterprise and retrieve the artifact. Data offers him some Klingon bloodwine that Worf replicated to his liking. Instead he simply looks at the beverage and pours it on the floor. Worf and Beverly are running out of time and excuses for the inspection, and have found nothing substantial. Before they can continue, the mercenary ship beams Riker and his team in the shuttle bay, with Riker stunning an operations crewman who was just entering. He explains to Beverly that he decided to make a career change with Picard demands Worf the location of the artifact, but Riker believes Koral has it and demands his location. Riker apologizes and stuns both Worf and Beverly; the latter seems to startle Picard a bit. While Data and Troi continue to fail at their attempts at conversation with Koral, the mercenary team beams in, which surprises all three of them. Picard retrieves the artifact from Koral and prepares to leave. Data warns Riker that his actions would bring charges of theft, assault, piracy and treason on him. Riker shrugs it off and decides to add one more to the list: He fires at Picard who rolls out of the way, and returns fire, sending Riker to the ground. Troi and Data rush to his aid and tells Picard that Riker is dead, which Picard takes pleasure of. They then return to the mercenary ship, as Riker comes to as he was stunned, which Data feels a similar sensation. Dazed, Riker explains it will take time to explain everything. Back on the mercenary ship, Picard delivers two items to Baran. The artifact, and a vicious punch to the face. This forces Picard to incite a mutiny on the ship, which he succeeds as the crew turns their back on Baran and follows Picard. Picard then explains that he switched the transponder codes. He destroys the control device and takes command. Next he orders the crew to deliver the artifacts as planned. On the Enterprise, Riker notifies Vulcan security minister Satok that their operative is on her way, and learns from the minister that the Vulcans have no one stationed on any mercenary ship. Meanwhile, Picard talks with Tallera about the two pieces, taking note of the symbol of War and Death and notices that the there seems to be an important glyph missing from between them. Tallera then says when they arrive that she will take the pieces and go to collect the payment. She says she must go alone, since she is Vulcan, to avoid suspicion. Picard then reveals that he asked Riker to contact Vulcan security, which unnerves Tallera. When they arrive at Vulcan , Picard stops her as she prepares to go, ordering her to leave one piece behind. Tallera then reveals to the rest of the crew his true identity as a Starfleet.

DOWNLOAD PDF A TRAITOR TO MEMORY, (2 PART SET: PART 1 12 TAPES PART 2 9 TAPES)

Chapter 7 : Gambit, Part II (episode) | Memory Alpha | FANDOM powered by Wikia

As I said, I do not think the case can be made for GWB as a traitor, and so answered in quote #1. Then you pointed out an inconsistency that simply does not exist in post #2, saying " This poster (Lemon Law) reveals an interesting posting technique you use to obscure the truth and provide the appearance of a well thought out case.".

This section is organized as follows: Overview of Memory Hierarchies 6. Basics of Cache and Virtual Memory 6. Beginning with latches, relays, and flip-flops, storage progressed to mercury delay lines an early form of dynamic memory , core memories comprised of magnets threaded on a grid of conducting wires, magnetic drum and disk storage, optical and magneto-optical storage, and faster memory, cache, and registers. It is natural to group these different types of memory, in particular the memory technologies used today, in terms of storage capacity and speed. This leads us to formulate an abstraction called a memory hierarchy, which we discuss in Section 6. Because certain types of memory are faster than others, and since designers place these faster types of memory closer to the processor, it is useful to put smaller amounts of data in the faster memory, which are subsets of larger data stored in slower memory. We call this technique caching pronounced "cashing" , which is discussed in Section 6. The analysis of memory performance is discussed in Section 6. We complete the discussion of the five components of a computer with an overview of input and output techniques, protocols, and devices e. Information contained herein was compiled from a variety of text- and Web-based sources, is intended as a teaching aid only to be used in conjunction with the required text , and is not to be used for any commercial purpose. Particular thanks is given to Dr. Enrique Mafla for his permission to use selected illustrations from his course notes in these Web pages. Registers are the fastest type of memory, which are located internal to a processor. These elements are primarily used for temporary storage of operands, small partitions of memory, etc. Cache is a very fast type of memory that can be external or internal to a given processor. Cache is used for temporary storage of blocks of data obtained from main memory read operation or created by the processor and eventually written to main memory write operation. Main Memory is modelled as a large, linear array of storage elements that is partitioned into static and dynamic storage, as discussed in Section 2 of these notes. Main memory is used primarily for storage of data that a program produces during its execution, as well as for instruction storage. Disk Storage is much slower than main memory, but also has much higher capacity than the preceding three types of memory. Because of the relatively long search times, we prefer not to find data primarily in disk storage, but to page the disk data into main memory, where it can be searched much faster. Archival Storage is offline storage such as a CD-ROM jukebox or in former years rooms filled with racks containing magnetic tapes. This type of storage has a very long access time, in comparison with disk storage, and is also designed to be much less volatile than disk data. The different partitions of the memory hierarchy each have characteristic persistence volatility. For example, data in registers typically is retained for a period of several cycles to hundreds of cycles, whereas data in cache can persist over tens to thousands of cycles. Data in main memory often remains for as long as the life of a program, and disk storage has persistence ranging from a few milliseconds to years if the disk is properly maintained. Archival storage lifetime can vary from hours e. The size of the different memory hierarchy partitions is also an important implementational consideration. For example, register files vary from a few hundred bytes to a few kilobytes, and modern caches from 64KB embedded processors to 2MB large engineering workstations. Main memory ranges from 1MB embedded processors to MB desktop computers to tens of gigabytes supercomputers. Disk storage ranges from tens to thousands of gigabytes, the latter in the case of disk arrays, whereas archival storage can range as high as hundreds of terabytes for the storage of remote sensing image data. Access times for registers are on the order of one CPU clock cycle, similar to cache. Main memory takes more time, perhaps several clock cycles to tens of cycles, depending on the disparity between the memory and CPU bandwidths. Disk storage latencies are typically on the order of milliseconds, and archival storage with the exception of CD-ROM jukeboxes takes minutes to days to access stored data. Each of these parameters - usage, persistance, size, and access time - are important to the design of storage systems. In the following section, we discuss how properties or attributes of each partition in the memory hierarchy influences the movement of data back and forth between levels of the memory hierarchy. In particular we focus on caching i. Here, we use the term data to generalize all bit patterns whether instructions or data. Putting aside the difference between cache and VM i. We begin with an event that occurs when data cannot be found in the cache or VM page store, called a cache miss or a page fault. The cache contains blocks of data from memory, and the VM contains pages of data from disk. Apart from the fact that cache blocks are usually smaller than memory pages, the concept is the same for cache and VM. Assume that we want to access a block in cache via a memory operation O. Using an address translation mechanism described in Section 6. When the processor accesses C using O, it performs a comparison operation to see if the block is really in the cache. If the block is present, this is called a hit, otherwise we have a miss. When the processor accesses P using I, it performs a comparison operation to see if the page is really in the page store. If the page is present, this is called a page hit - otherwise, we have a page fault. In cache, if the block is not found, then it is retrieved from main memory and stored in the cache. At this point, memory operation O can proceed. Schematic diagram of a cache and b virtual memory, showing the similarities between the two techniques. Thus, cache and VM both have similar errors. These misses or page faults adversely impact memory performance at their respective levels of the memory hierarchy, via the following three mechanisms: Replacement cost - When a cache miss or VM page fault occurs, the processor or memory management unit must 1 find the requested block or page in the lower-level store, 2 write the block or page back to the cache or paging store, and 3 set the appropriate bits per Section 6. These operations incur overhead that would not have occurred if the memory operation had been successful i. Of particular interest are the bus and memory costs discussed below , as well as the block or page replacement strategy. It is crucial that page replacement cause the least possible number of misses or page faults, as discussed in Section 6. Bus cost - When a cache block is obtained from memory, it must be transmitted along the bus that connects main memory with the processor or cache depending on how the physical architecture supports cache. This bus must have very high bandwidth, so as not to impede cache replacement. Similarly, in VM, the bus between the disk and memory must be sufficiently fast not to inhibit data transfer in support of page replacement. Lower-level storage cost - If a cache miss is encountered, then the requested data must be retrieved from main memory, which implies that main memory itself must be sufficiently fast not to impede overall cache performance for a given average hit frequency. In order to make memory access efficient under the assumption of cache or virtual memory or both , we need to correctly design the buffering function implemented in cache or VM by making the computationally intensive parts of the buffering process as efficient as possible. In the next section, we show that the potentially costly operation of address translation can be made efficient via hardware implementation. Address Translation Mechanisms for Cache and VM One of the key problems of caching or virtual memory systems is how to translate the virtual address which the CPU uses into the corresponding physical address cache block or page address in the paging store. This can be challenging because of a the size of the virtual address space in virtual memory, and b temporal inconsistencies between cache and main memory contents, or between paging store and disk contents in VM. Let us derive an address translation scheme using cache as an example. The cache is a linear array of entries, each entry having the following structure: Data - The block of data from memory that is stored in a specific entry also called row or collection of entries in the cache, Tag - A small field of length K bits, used for comparison purposes in verifying correct addressing of data, and Valid Bit - A one-bit field that indicates status of data written into the cache. Assume that the N-bit virtual address the address that the processor uses to access cached data is divided into three fields: Tag - A K-bit field that corresponds to the K-bit tag field in each cache entry, Index - An M-bit field in the middle of the virtual address that points to one cache entry, and Byte Offset - Two bits in MIPS that are not used to address data in the cache. In order to understand how the cache entries are referenced by the index, we need to observe that there are three ways of arranging data in a cache: Direct Mapped - There is a one-to-one correspondence between each block of data in the cache and each memory block that we are trying to find. That is, if we are trying to find a memory block b, then there is one and only one place in the cache where b is stored. Set Associative - Given memory block b that we are trying to find in the cache, there are L entries in the cache that can contain b. We say that this type of cache is called L-way set associative. Fully Associative - Any block of memory that we are looking for in the cache can be found in any cache entry - this is the opposite of direct mapping. As shown in Figure 6. Here, bits are occupied by the Tag field, bits contain the Index field, and bits 0,1 contain the Byte Offset information. The index points to the row or entry in cache that supposedly contains the data requested by the processor. We say "supposedly" because the requested data block might or might not be in cache at the location pointed to by the Index. If the tags do not match, then a cache miss is detected and the comparator shown as a circle with an equal sign in Figure 6. Otherwise, the comparator outputs a one, which is and-ed with the valid bit in the cache row pointed to by the Index field of the cache address. If the valid bit is a one, then the Hit signal output from the and gate is a one, and the data in the cached block is sent to the processor. The Virtual Address shown in Figure 6. The Valid Bit is then checked - if zero, the page is not in memory, and the page offset must be loaded into the Page Table. The key to improvement of performance in VM is thus efficient transformation of the Virtual Address into the Physical Address. The performance of a VM address translation mechanism can be improved by the insertion of a Translation Lookaside Buffer, which exploits the principle of locality. That is, when an address translation is performed, it will probably be required again soon due to spatiotemporal locality of page references. As a result, modern VM systems use an extra piece of hardware that acts a cache for recent translations. This piece of hardware is the TLB, which holds only page table mappings. The TLB operates as follows: If a hit results, then the physical page number is used to form the address per Figure 6. If the processor is performing a memory write, the dirty bit discussed in Section 6. Control is then passed to Step 1 of this algorithm, where the next page reference is processed. If a miss occurs, determine whether it is a page fault or a TLB cache miss, and act accordingly: The page reference Step 1 is then tried again. If a page fault page not in memory , then this condition is handled by the VM portion of the operating system through a page fault exception produced by the CPU. This process can be performed either in hardware or software, and there is little performance difference because the basic operations are the same, as discussed below. Because the TLB is much smaller has fewer entries than the page table, the probability of TLB misses is much larger than the probability of page faults.

DOWNLOAD PDF A TRAITOR TO MEMORY, (2 PART SET: PART 1 12 TAPES PART 2 9 TAPES)

Chapter 8 : Organization of Computer Systems: Processor & Datapath

**Traitor in Zebra* [, Blackman] -- (The title doesn't show up on the case or the disc but it does appear on the disc menu - it's supposed to be on Volume 4.) A British Naval radar tracking system is being jammed by enemy agents.*

Overview[edit] Non-volatile memory is typically used for the task of secondary storage , or long-term persistent storage. However, most forms of non-volatile memory have limitations that make them unsuitable for use as primary storage. Typically, non-volatile memory costs more, provides lower performance, or have limited lifetime compared to volatile random access memory. Non-volatile data storage can be categorized into electrically addressed systems read-only memory and mechanically addressed systems hard disks , optical disc , magnetic tape , holographic memory , and such. Non-volatile random-access memory Electrically addressed semiconductor non-volatile memories can be categorized according to their write mechanism. Programmable read-only memory can be altered after manufacture, but require a special programmer and usually cannot be programmed while in the target system. The programming is permanent and further changes require replacement of the device. Data is stored by physically altering burning storage sites in the device. EPROMs have a quartz window that allows them to be erased with ultraviolet light, but the whole device is cleared at one time. These erasable memory devices require a significant amount of time to erase data and to write new data; they are not usually configured to be programmed by the processor of the target system. Data is stored by use of floating-gate transistors which require special operating voltages to be applied to trap or release electric charge on an insulated control gate for storage sites. Flash memory The flash memory chip is a close relative to the EEPROM; it differs in that it can only erase one block or "page" at a time. It is a solid-state chip that maintains stored data without any external power source. NOR flash provides high-speed random access, reading and writing data in specific memory locations; it can retrieve as little as a single byte. NAND flash reads and writes sequentially at high speed, handling data in small blocks called pages, however it is slower on read when compared to NOR. NAND flash reads faster than it writes, quickly transferring whole pages of data. Magnetoresistive random-access memory Magnetoresistive RAM is one of the newest approaches to non-volatile memory and stores data in magnetic storage elements called magnetic tunnel junctions MTJs. The second generation is developed mainly through two approaches: IBM Millipede and holographic memory Mechanically addressed systems utilize a contact structure "head" to read and write on a designated storage medium. Since circuitry layout is not a key factor for data density, the amount of storage is typically much larger than for electrically addressed systems. Tape media can be removed from the drive and stored, giving indefinite capacity at the cost of the time required to retrieve a dismounted tape. Formerly, removable disk packs were common, allowing storage capacity to be expanded. Optical discs store data by altering a pigment layer on a plastic disk, and are similarly random access. Read-only and read-write versions are available; removable media again allows indefinite expansion, and some automated systems were used to retrieve and mount disks under direct program control.

Chapter 9 : ShieldSquare reCAPTCHA Page

Sue: Robert racedaydvl.com God for that. I wonder which one of them will turn out to be the traitor My money's on Vila. On the planet Helotrix, a Federation Colonel named Quute is interrogating an insurrectionist named Igin.