

Chapter 1 : Illustrated GOF Design Patterns in C# Part VI: Behavioral III - CodeProject

Part V Lattices and Geometry Lattices and Cryptography I. F. Blake A Simple Construction for the Barnes-Wall Lattices G. Nebe E. M. Rains N. J. A. Sloane Part VI Behaviors and Codes on Graphs.

We will discuss the Strategy, Template Method, and Visitor patterns. I will also provide some closing remarks. It is assumed the reader is familiar with basic C syntax and conventions, and not necessarily details of the. This article will only give a brief overview of the pattern, and an illustration of the pattern in C. A design pattern is not code, per se, but a "plan of attack" for solving a common software development problem. The GOF had distilled the design patterns in their book into three main subject areas: Creational, Structural, and Behavioral. This article deals with the Behavioral design patterns, or how objects act. The first three articles in this series dealt with the Creational and Structural patterns, and the last article continued the Behavioral patterns. This article concludes the illustration of the Behavioral patterns as described by the Gang of Four. This article is meant to illustrate the design patterns as a supplement to their material. It is recommended that you are familiar with the various terms and object diagram methods used to describe the design patterns as used by the GOF. The most important terms to get your head around are abstract and concrete. The former is a description and not an implementation, while the latter is the actual implementation. In C, this means an abstract class is an interface, and the concrete class implements that interface. Behavioral Patterns To quote the GOF, "Behavioral patterns are concerned with algorithms and the assignment of responsibilities between objects. Behavioral patterns describe not just patterns of objects or classes but also the patterns of communication between them. They shift your focus away from flow of control to let you concentrate just on the way objects are interconnected. For example, a text-rendering engine may lay out text as plaintext in one case, rich text in another, and as cells in a spreadsheet in yet another. While the activity is the same, composition, the algorithm used varies. The Strategy behavioral design pattern is useful in that it "define[s] a family of algorithms, encapsulate[s] each one, and make[s] them interchangeable. Strategy lets the algorithm vary independently from the clients that use it. The GOF suggest use of this pattern when many related classes differ only in their behavior. Strategies provide a way to configure a class with one of many behaviors. Strategies can be used when these variants are implemented as a class hierarchy of algorithms. Use the Strategy pattern to avoid exposing complex, algorithm-specific data structures. Instead of many conditionals, move related conditional branches into their own Strategy class. The structure of this pattern appears thus: For our simple example strategy. We create the abstract strategy: MakeList "This is a sample of varying the strategy" ; lc. SetStrategy new SortedStrategy ; lc. Output ; Although this is a simple example, one can see that the Strategy pattern allows great flexibility in algorithm implementation and greatly simplifies object activity. We can "plug in" additional strategies quite easily. Template Method The Template Method pattern allows software to define a skeleton of an algorithm in an operation, deferring some steps to subclasses. We see this all the time in virtual functions in C, therefore I shall not illustrate this pattern with a code example, as we have been using Template Methods throughout these articles. The Template Method structure: Visitor The Visitor pattern "Represent[s] an operation to be performed on the elements of an object structure. Visitor lets you define a new operation without changing the classes of the elements on which it operates. It is possible using this pattern to separate the structure of an object, from the operations that act on the structure. This is useful in that new operations may be defined without modifying the original object. As a trivial example, a bank account object may be operated on for credit or debit purposes. Later, it is possible to add an "interest" operation, a "service fee" operation, or "transmit information to mobile device" operation without modifying the original object. The GOF suggests using the Visitor pattern when an object structure contains many classes of objects with differing interfaces, and you want to perform operations on these objects that depend on their concrete classes. Visitor lets you keep related operations together by defining them in one class. When the object structure is shared by many applications, use Visitor to put operations in just those applications that need them. Changing the object structure classes requires redefining the interface to all visitors, which is potentially costly. The Visitor pattern structure: For our example visitor. We will simplify

the example by not taking any date-related information into account. Products are the elements that need to be "visited. Our client would then iterate through the products, using the appropriate Visitor: Add new Product "chair", 5, 2, Add new Product "desk", 10, 5, Add new Product "filing cabinet", 20, 7, Accept cv ; p. Accept sv ; p. Results ; It would be fairly simple to add a new type of BaseSummingVisitor to our application, for example a TaxVisitor. This is a powerful design pattern, prone to misuse if not properly thought out. I encourage you to review its nuances in Design Patterns. Conclusions Behavioral design patterns allow for great flexibility in how your software behaves: The Strategy pattern is useful because it allows great flexibility in client usage of a family of algorithms to accomplish a task. You can separate the data used in implementation of an algorithm from clients that does not need to know about, thereby reducing complexity, class bloat and areas of highly branched logic that would otherwise exist. One must recognize some drawbacks as well: Template Methods allow us to GOF: You can define a template method that calls "hook" operations at specific points, thereby permitting extensions only at those points. The Visitor pattern allows separation of the structure of objects from the operations that act on them. You can add operations easily, gather related operations together and separate unrelated ones. There are some limitations: Adding new elements to be operated on can be difficult, because they may add new abstract operations which need to be implemented in each concrete visitor. One may also break encapsulation because of the need to operate on data contained within the visited elements. Please review the Design Patterns Visitor consequences for more information. They, as well as other design patterns, should not be treated as a complete solution for software development, rather as a set of guiding principles. Good programming is still paramount - without a good programming foundation, all the patterns in the world are useless. There are times to break from known patterns, to strike out on your own, leading you to new, innovative patterns or solutions to software design and implementation. Recognize and expect limitations while also recognizing opportunity Armed with good programming and design skills, it should truly be a golden age. Stay tuned for future articles Building the Samples Unzip the source files to the folder of your choice. Start a shell cmd. You may have to alter the Makefile to point to the correct folder where your. NET Framework libraries exist. Addison Wesley Longman, Inc.

Chapter 2 : Data and Science Series Archives - George Dell, SRA, MAI, ASA, CRE

Codes, Graphs, and Systems is an excellent reference for both academic researchers and professional engineers working in the fields of communications and signal processing. A collection of contributions from world-renowned experts in coding theory, information theory, and signal processing, the book provides a broad perspective on contemporary.

There are lots of beautiful distractions which are competing for my attention, including my gorgeous 20 month old granddaughter who is playing at my feet. So I have to disclaim any liability related to the quality of this particular blog! So far in this series we have discussed how to set the coefficients for the PI controllers in a cascaded speed control loop. We saw that K_b is used for pole-zero cancellation within the current loop. But what should that bandwidth be? It turns out that we have two competing effects vying for control of where we set K_a . On the bottom end of the frequency range, we have the velocity feedback filter pole with some really nice tan lines that wants to push the current controller pole to higher frequencies so as to not interfere with our nice tuning procedure. But we can only go so high before we start running into other problems. Figure 1 shows an example case to illustrate this point. The green curve represents what the tuning procedure predicts the normalized step response should be for a system with a damping factor of 2. The system is still stable, but the damping is much less than predicted from our tuning procedure. At this point, we have two options—either increase the damping factor and consequently lower the frequency response of the velocity loop, or increase the current loop bandwidth by increasing K_a . The cyan curve shows the first option where we increase the damping factor just enough to bring the overshoot down to the predicted value. Unfortunately this increases the step response transient time as well. The yellow curve shows the latter option where we put the damping factor back to 2. As you can see, the actual response more closely resembles the predicted value.

Normalized Small Signal Step Response

So from this exercise, we might conclude that the best strategy is to set the current controller bandwidth to be as high as possible. But is this really the best course of action? Usually high bandwidth in any system results in unruly and obnoxious behavior, and should only be used if absolutely necessary. In this case, high current loop bandwidth often results in undue stress on your motor, since high frequency current transients and noise translate into high frequency torque transients and noise. This can even manifest itself as audible noise! But there is also another limit on your current loop bandwidth: To simplify the discussion, we will assume that the entire control loop is clocked by a common sampling signal, although in real-world applications we often choose to sample the velocity loop at a much lower frequency than the current loop to save processor bandwidth. In an analog system, any change in the motor feedback signals immediately starts having an effect on the output control voltages. But with the digital control system of Figure 2, the motor signals are sampled via the ADC at the beginning of the PWM cycle, the control calculations are performed, and the resulting control voltages are deposited into double-buffered PWM registers. From a system modeling perspective, this looks like a sample-and-hold function with a sampling frequency equal to the PWM update rate frequency. The fixed time delay from the sample-and-hold shows up as a lagging phase angle which gets progressively worse at higher frequencies. Figure 3 shows a normalized frequency plot of the phase delay for a sample-and-hold function, where the sampling frequency is assumed to be 1. As you can see, the phase delay reaches down into frequencies much lower than the sampling frequency. However, for a digital control system, you must add the phase lag shown in Figure 3. To do this, calculate for the following frequency ratio: L is the motor inductance t is the time constant of the velocity filter d is the damping factor T_s is the sampling period In most designs this still leaves a fairly broad range for the value of K_a . In my experience, I have found that erring on the side of the current bandwidth being too low is usually better than being too high. Snappy waveforms look nice on an oscilloscope, but your application may not need or want all that bandwidth. In fact, it will usually just make the overshoot worse, since the integrator is acting on a gained-up error signal, which it will just have to dump eventually. So how do we deal with this problem? Are we doomed to simply using low integrator gains? Until then—Keep Those Motors Spinning, Ros River

Thank you very much for the perfect bridge between highschool theory and the real industrial application.

DOWNLOAD PDF PART VI BEHAVIORS AND CODES ON GRAPHS

Could you please answer two question regarding the provided Excel-Spreadsheet: What happens to that value in an application with an encoder?

Chapter 3 : Teaching Your PI Controller to Behave (Part VI) - Motor Drive & Control - Blogs - TI E2E Comm

Using the example of the cognitively impaired resident in the psychiatric facility who has ongoing exit behaviors, think in terms of how to remove the risk for this and other individuals. Recall that this behavior was known to staff and was suggested.

Remote Control this article In the past two articles I have implemented the motor and distance sensor control for my little robot vehicle, using sound software development techniques that make this project easier to write and to maintain. When I started writing the remote control code I became a bit frustrated with having to constantly connect and disconnect Bluetooth and USB. To avoid this hassle I decided to use the SoftwareSerial. The change is pretty simple. I moved the connections from the Bluetooth module going into pins 0 and 1 to analog pins 2 and 3 numbered 16 and 17 , which are unused. Then I created a SoftwareSerial object at the top of the sketch: The setup in my sketch then changed to: Because I elected to keep sending the logging output to the regular serial interface. I now have access to two serial ports, so I will use the Bluetooth connection for remote commands and the USB one for logging and debugging output. With the Bluetooth module connected in this way there is no collision with the USB serial interface so I could keep the USB cable connected at all times while developing and testing. But this case is a little different. We really have two separate pieces to abstract. The most obvious is the communication method. My robot uses Bluetooth, but this is not the only mechanism by which a robot can be controlled. I clearly need a device driver that can read remote commands, regardless of the wireless technology used. The second, and less obvious dimension to this problem is the remote protocol. Different remote controllers will use different command codes to mean the same thing. For example, one remote control may send the command to turn left as a byte, maybe the letter "L". Another controller with a numeric keypad configuration may use "4" for the same thing. Another, joystick inspired remote control may not support a direct "go left" command at all and may just send instructions as vectors in a two-dimensional grid, so going left would be given as two numbers, -1, 0. I do not really want to deal with these differences in the high-level robot code because that will make my code more complex. So clearly a driver that can hide these protocol differences from the main sketch is also needed. So do I need two different drivers, then? Ideally I want something simple, like a getCommand function that just returns the next command if one is available. There are two ways this could work. This provides great flexibility, because the same protocol can easily be supported over different communication channels. And what would be the alternative? I could embed the communication part inside the protocol driver. This is less flexible because now the remote protocol is tied to a particular communication channel. Which of these options is better? If I was working with standard protocols that remote controls of different kinds implement, then having separate protocol and communication drivers makes sense. In this case, however, I believe the kinds of remote controls that I will find are mostly using proprietary protocols. My impression is that I will never need to use a protocol driver with more than one communication method. I can always break the driver into two in the future if I find the need. The remote control driver interface Given the variety of remote controls out there, I need to find a common format that all the remote drivers can use to represent commands. For the purpose of driving a vehicle, I can think of three types of remote controls that I might ever need to support: In all cases it is safe to assume that in addition to the directional control there are a few function buttons. As you can see there are significant differences between the types of remote controls, so the job of defining a common format for all of them is not easy. Obviously I cannot indicate direction with simple forward, backward, left and right commands, since many remote controls can provide a finer degree of directional control and it would be a pity to ignore all that extra information. So I have to pick one of the more advanced remote control data representations. The one that translates more easily into motor instructions for my vehicle is the one where the controller uses two sliders, each controlling one side of the vehicle. This directly translates into speeds for my two motors, so the application can simply send the data from the remote control into the motors. But there is a downside to selecting a single data representation for all types of controllers. To help with this task the base driver class will provide the conversion functions for all types of

remote controls. The representation of a command includes left and right slider values and a possible function key. A remote control driver can provide up to four function keys, all listed in the enum definition. The `getRemoteCommand` method will be implemented by remote control drivers. The first five just set the left and right sides appropriately, while the sixth copies the values given as arguments, since these match the internal data representation. The next conversion is the one that requires some thought. Here I have a controller that has two sliders, one vertical to move forward or backward and one horizontal to move left or right. Somehow I need to use these two values to derive independent left and right values to power the motors. The algorithm that I came up with for this conversion is relatively simple. And these changes will just make the robot go towards the left side, which is what I want. The last conversion function is for joystick type controllers.

Remote control driver implementation My choice of controller is a free Android app called BlueStick. Here is a screenshot of this app: This is a very simple controller with five directional commands and six function buttons. The directional commands can be triggered by touching arrow buttons on the screen or by tilting the phone. The app documentation provides the codes that are sent over Bluetooth for each of the commands: The above list defines the protocol for this remote control, so this is really all I need to know to be able to implement this driver. So here is the code for the BlueStick driver: This implementation is extremely simple. If the `BTSerial` class has a character in its queue, then I read it, and depending on what character it is I configure the command, using conversion functions for the directional commands or the key constants for the function keys. I have six function keys in this remote control app, but I decided to just use four in my driver, so I map the two extra keys as duplicates of another two. I also need to explain what the `lastKey` member variable is for. As a test I connected the BlueStick app running on my cell phone to a Bluetooth terminal running on another phone, just to confirm that the commands the app sends are in alignment with the documentation. I found that they did match, but also found that the app is constantly sending commands. If I find that the new key is the same as the previous one then I just throw it away as a duplicate. Another important thing to note about this implementation is that it tries to be tolerant of unknown codes. At the start I initialize the command structure as a stop with no function keys. If I receive an unknown character then the sketch will receive a stop command, which will make the robot stop and wait for more commands. My goal for the remote controlled robot is to incorporate the automated mode I wrote in the previous article as an option that can be enabled with a special remote command. After some consideration I designed the following flow chart for how the robot will operate under remote control: The robot will begin by listening for remote commands, without moving. If a command is received it will be executed, and then it will go back to listen for more commands. There will be a "Roomba" mode command that acts as a toggle. Any other commands sent while the robot is in "Roomba" mode will be ignored. Adding remote control to the sketch Okay, now I will move on to the most exciting part, which is to actually write the code to make the robot follow the remote commands. Starting from the sketch as I left it in the previous article I begin by adding the remote control driver class to the top of the sketch, where all the drivers are: Now I can add a remote control object to the Robot class: I now need to have an additional state for when the robot is under remote control: Again, to be consistent I need one for the new state: In case I need to do debugging at some point I kept the log statement I had before, enhanced to include remote control data as well. For the remote control I have to note if I have received a command or not. This is what `haveRemoteCmd` does. If this variable is true then the `remoteCmd` struct holds the command data. If the variable is false then `remoteCmd` has an undefined value and should not be used. Once I gathered all my inputs I move on to control the robot behavior. In this new version of the sketch there are two different operating modes. The `if remoteControlled` statement takes care of differentiating between the two modes. The `if` portion deals with the remote control mode, while the `else` portion does the "Roomba" mode. If there is a remote command available I then look at the key value that came with it. If the key is `keyF1` value then I enable "Roomba" mode by calling `move` see the previous article for an explanation of this method. If the command came without a key `keyNone` then it is a directional command, so all I do is move the left and right values from the remote command into the motors. Believe it or not this short piece of code fully handles the remote control of the vehicle! The "Roomba" mode handling is mostly the same as in the previous article, with the exception that I removed the 30 second timeout handling that I had in that

version and replaced it with this: The rest of this part deals with the "Roomba" mode logic, and this did not change from the previous version of the sketch. My choice is RocketBot , another Android app. Here is a screenshot of RocketBot:

Chapter 4 : Table of contents for Library of Congress control number

The graph_iter_concat() answers a question I had about how graph_iter_t is used, here is how it looks like: This is using C's support for passing an unknown number of parameters to a method. This basically builds a simple linked list, which also answers the usage of _blarf() function and its behavior a few posts ago.

Introduction Deep reinforcement learning is on the rise. No other sub-field of Deep Learning was more talked about in the recent years - by the researchers as well as the mass media worldwide. Most outstanding achievements in deep learning were made due to deep reinforcement learning. AI agent learned how to run and overcome obstacles. Other AI agents exceed since human level performances in playing old school Atari games such as Breakthrough Fig. The most amazing thing about all of this in my opinion is the fact that none of those AI agents were explicitly programmed or taught by humans how to solve those tasks. They learned it by themselves by the power of deep learning and reinforcement learning. The goal of this first article of the multi-part series is to provide you with necessary mathematical foundation to tackle the most promising areas in this sub-field of AI in the upcoming articles. The environment may be the real world, a computer game, a simulation or even a board game, like Go or chess. Like a human the AI Agent learns from consequences of its Actions, rather than from being explicitly taught. The neural network interacts directly with the environment. It observes the current State of the Environment and decides which Action to take e. The amount of the Reward determines the quality of the taken Action with regards to solving the given problem e. The objective of an Agent is to learn taking Actions in any given circumstances that maximize the accumulated Reward over time. MDP is the best approach we have so far to model the complex environment of an AI agent. The agent takes actions and moves from one state to an other. In the following you will learn the mathematics that determine which action the agent must take in any given situation. This is also called the Markov Property Eq. For reinforcement learning it means that the next state of an AI agent only depends on the last state and not all the previous states before. In a Markov Process an agent that is told to go left would go left only with a certain probability of e. With a small probability it is up to the environment to decide where the agent will end up. S is a finite set of states. P is a state transition probability matrix. Here R is the reward that the agent expects to receive in the state s Eq. This process is motivated by the fact that for an AI agent that aims to achieve a certain goal e. The primary topic of interest is the total reward Gt Eq. It is mathematically convenient to discount rewards since it avoids infinite returns in cyclic Markov processes. Besides the discount factor means the more we are in the future the less important the rewards become, because the future is often uncertain. If the reward is financial, immediate rewards may earn more interest than delayed rewards. The value function maps a value to each state s. The value of a state s is defined as the expected total reward the AI agent will receive if it starts its progress in the state s Eq. The value function can be decomposed into two parts: This function can be visualized in a node graph Fig. Starting in state s leads to the value v s. In this particular case we have two possible next states. Thus the immediate reward from being in state s now also depends on the action a the agent takes in this state Eq. Mathematically speaking a policy is a distribution over all actions given a state s. The policy determines the mapping from a state s to the action a that must be taken by the agent. The policy leads to a new definition of the the state-value function v s Eq. Notice that for a state s, q s,a can take several values since there can be several actions the agent can take in a state s. The calculation of Q s, a is achieved by a neural network. Given a state s as input the network calculates the quality for each possible action in this state as a scalar Fig. Higher quality means a better action with regards to the given objective. Action-value function tells us how good is it to take a particular action in a particular state. Previously the state-value function v s could be decomposed into the following form: The same decomposition can be applied to the action-value function: At this point lets discuss how v s and q s,a relate to each other. The relation between these functions can be visualized again in a graph: In this example being in the state s allows us to take two possible actions a. By definition taking a particular action in a particular state gives us the action-value q s,a. Now lets consider the opposite case in Fig. The root of the binary tree is now a state in which we choose to take an particular action a. Remember that the Markov Processes are stochastic. Strictly

speaking you must consider probabilities to end up in other states after taking the action. Now that we know the relation between those function we can insert v s from Eq. This recursive relation can be again visualized in a binary tree Fig. To obtain q s , a we must go up in the tree and integrate over all probabilities as it can be seen in Eq. Furthermore the agent can decide upon the quality which action must be taken. The best possible action-value function is the one that follows the policy that maximizes the action-values: To find the best possible policy we must maximize over q s , a . Maximization means that we select only the action a from all possible actions for which q s , a has the highest value. Thus provides us with the Bellman Optimality Equation: The agent knows in any given state or situation the quality of any possible action with regards to the objective and can behave accordingly. Solving the Bellman Optimality Equation will be the topic of the upcoming articles. In the following article I will present you the first technique to solve the equation called Deep Q-Learning. It has the emphasis on building Deep Learning Applications in the field of Predictive Analytics and making it work in a Production Environment.

5 Part VI. Factor the polynomials completely when given 1 factor of the polynomial. Factor $3^4 \cdot 2^8 \cdot 16x \cdot x^3 \cdot 2$ completely given $(2)x$ is a factor. Factor $3 \cdot 16 \cdot 3 \cdot 10x \cdot x$

This book provides practical guide to cluster analysis, elegant visualization and interpretation. It contains 5 parts. Part I provides a quick introduction to R and presents required R packages, as well as, data formats and dissimilarity measures for cluster analysis and visualization. Part II covers partitioning clustering methods, which subdivide the data sets into a set of k groups, where k is the number of groups pre-specified by the analyst. Partitioning clustering approaches include: In Part III, we consider hierarchical clustering method, which is an alternative approach to partitioning clustering. The result of hierarchical clustering is a tree-based representation of the objects called dendrogram. In this part, we describe how to compute, visualize, interpret and compare dendrograms. Part IV describes clustering validation and evaluation strategies, which consists of measuring the goodness of clustering results. Among the chapters covered here, there are: Assessing clustering tendency, Determining the optimal number of clusters, Cluster validation statistics, Choosing the best clustering algorithms and Computing p-value for hierarchical clustering. Part V presents advanced clustering methods, including: Hierarchical k-means clustering, Fuzzy clustering, Model-based clustering and Density-based clustering. He works since many years on genomic data analysis and visualization read more: He has work experiences in statistical and computational methods to identify prognostic and predictive biomarker signatures through integrative analysis of large-scale genomic and clinical data sets. He created a bioinformatics web-tool named GenomicScape www. He is the author of many popular R packages for: Recently, he published three books on data analysis and visualization: Practical Guide to Cluster Analysis in R <https>: Complete Guide to 3D Plots in R <https>:

In my case, I was using the FAST observer which is part of our InstaSPIN-FOC product. With InstaSPIN-FOC, the filter value is selectable. For my system, I was able to get satisfactory performance with a pole value of rad/sec², but admittedly, this was a subjective decision.

We know how LemonGraph starts to execute a query. It find the clause that is likely to have the least number of items and starts from there. These are the seeds of the query, but how is it going to actually process that? Here is my graph: And here is the query that I want to run on it: The actual behavior starts from the matches method, which starts here: As usual for this codebase, the use of tuples for controlling behavior is prevalent and annoying. In this case, the idx argument controls the position of the target in the query, whatever this is the start, end or somewhere in the middle. The deltas define what direction the matches should go and the stops where you must stop searching, I think. Now that we setup the ground rules, the execution is as follows: In this case, because the seed of our query is the edge which is in the middle it determines that deltas are 1, -1 and stops are 2, 0. The link variable there controls whatever we should follow links going in or out. There is a lot of logic actually packed into the link initialization. In this case, we are being passed: As you can see, we first validate that the match is valid this is where we mostly check other properties of the value that we are currently checking, filtering them in place. The first line there shows usage of delta, which control in what direction to move. This looks like redundant code that snuck in somehow. On the other hand, the Node. In that case, that would give a performance boost because it could filter a lot of edges in busy graphs. It took me a few reads to figure out that this is probably a case of someone unpacking a statement for debugging but forgetting to remove the packed statement. The second return statement is ignored and likely stripped out as unreachable, but it is pretty confusing to read. Now I know a lot more about the behavior of the code, so this make sense. These update the chain, like so: In processing the query, we first append the edge to the chain: Remember that in the matches , we got into this if clause: Look at the push method there. By the time we get to the result , we have iterated both sides of the connection, giving us: That is a very clever way of doing things. That will allow me to figure out the other side of this operation. I just run this through the debugger, and this is indeed how it actually works. The usage of yield to pass control midway is not trivial to follow, but it ends up being quite a nice implementation. This is enough for this post. In my next one, I want to explore a few of the possible operations that exposed by LemonGraph.

Chapter 7 : Part VI – CMS Immediate Jeopardy Series | Courtemanche & Associates

For example, the behaviors could occur multiple times during a given interval, but it is simply coded as one occurrence. Behaviors Investigated Chart #11 B is used to investigate multiple high frequency behaviors.

Event to command Part VI: Behaviors for coping with orientation changes Part VII: Your ViewModel and Model code can be placed in a portable class library, thus making it reusable. But, since it is not possible to put any platform specific code in your PCL, you cannot directly use platform specific services from within a ViewModel this would make your PCL unsharable between platforms. This poses some challenges, but these can quite easily be overcome. This part in our series will be the most extensive thus far, and will contain the most advanced concepts, like inversion of control, dependency injection, and reflection. But bear with me, it will all be worth it in the end! Windows Phone and Windows 8 both allow you to navigate from one page to another. The way they do this, however, is different between the two platforms. They both have a Navigate method for this, but the method exists on two different types, and also, the actual parameterlist of this method is different on each platform. You will need to find a solution to get a comparable way of navigating between pages on both platforms. You can do this by wrapping the platform specific functionality with an interface. Your ViewModels can then talk to this interface instead of talking directly to the actual navigation service of each platform. I have for instance a MainViewModel which acts as start page for my application. From this MainViewModel you can navigate to other pages in the application. Since I will be navigating in my ViewModels and not in my Views, I decided that when I navigate to another page, this is actually the same as navigating to another ViewModel. Also, I want to be able, when I navigate, to send along an additional parameter, which contains extra data for the ViewModel I navigate to. In the above code you can see it is quite easy to navigate to another ViewModel, and send along extra data navigate to the SudokuBoardViewModel with a specific game level. But you can also choose not to use this extra data navigate to the SudokuRulesViewModel. The interface itself has one extra method for navigating back. What you will still have to do now is write platform specific implementations of the INavigate service for each platform you want to support. The good thing is, you will need to write this kind of implementations only once. After your first app, you can reuse this effort to a next application and I will also make sure my own little API will become available on github in the next couple of weeks. To navigate from one page to the next, you need to make use of your Window. Content property or rootframe. In a clean Windows Store project you get a reference to this in your OnLaunched event I set up the entire framework from the OnLaunched handler in the app. Next, we will also need some means of associating the ViewModel we wish to navigate to, with a certain View, since navigating in Windows Store applications is done to views and not to viewmodels. This viewlocator is nothing more than a dictionary that associates views to viewmodels. As you can see, I ask the viewlocator for the view that is associated with the viewmodel I wish to navigate to. Next I tell the rootframe to actually navigate to this view. The second parameter in the second line of the NavigateTo method instantiates my viewmodel. The inversion of control container has been set up with a list of all my viewmodels. The main advantage using and IoC container will be that any dependencies your viewmodel uses, will get injected by the IoC, as long as you registered them at startup. Navigating with the extra data is a bit more complex. I created a IHandle interface which indicates whether a ViewModel can handle a certain message. If it does, it implements the IHandle interface. Before navigating to the view, I first let the ViewModel handle the actual request. One more thing missing after navigating to a certain view, is the actual databinding between View and ViewModel. For this I created an extra base class for each View. This allows us to navigate from one ViewModel to another, while sending along data. Also we are now able to actually databind our ViewModel to the actual View. In Windows Phone you navigate based on a Uri. Also, we will not be able to send along our ViewModel and request as easily. In this case, we will serialize everything before we navigate. Again, we will use a MvvmPage base class, which all our views can inherit from. In its NavigatedTo event handler we will now have to re create our ViewModel and request. I pulled out this code to another class which implements ICreateViewModels. Since views cannot use dependency injection, I wrapped the IoC in a singleton, which can give me my

viewmodelcreator. The ViewModelCreator then uses a deserializer and some reflection to recreate the viewmodel and to make it handle the actual request. The above code is actually the most complex part of the entire framework. An, as said, code like this, you will only have to provide once, after that you can reuse it for your next applications. Once you got thi plumbing in place, you can navigate between ViewModels in a way that is reusable on Windows Store and Windows Phone applications. It wil make it easier to put your logic code in a shared library and to easily make changes to it that will get reflected on each platform. Hang on for the next part, where I will talk about some extra helper classes to make MVVM work better for you. Other parts in this series:

Part VI: Remote Control (this article) In the past two articles I have implemented the motor and distance sensor control for my little robot vehicle, using sound software development techniques that make this project easier to write and to maintain.

The survey team leader wants to address an Immediate Jeopardy situation with your leadership. The surveyor verbally presents the circumstances of the situation, the investigation performed and the decision process. The specific details, including the individuals at risk is provided before leaving the entity. What should you do? Immediately remove the risk! To take appropriate actions to remove the risk, you may need to perform your own investigation to assure you know what the situation is, how isolated or widespread. Using the example of the cognitively impaired resident in the psychiatric facility who has ongoing exit behaviors, think in terms of how to remove the risk for this and other individuals. Recall that this behavior was known to staff and 1: The lock was not secure, and the resident knew how to access the keypad " address the lock and keypad issue. The entity had not investigated the issue. The facility was not secure, and the resident was able to exit the building to a busy street, can this be addressed to prevent exit, secure the premises, monitor individuals activity? Think not only of how to protect this individual but all individuals who may be affected by these circumstances. Address the who, what, when, where and why questions. Can it be addressed? Was exit behavior addressed in the care plan? Were actions carried out? Is the plan effective? How many other residents are at risk? How did the resident elope? What supervision was provided? Was outside medical treatment needed? When was the resident last seen? When the resident was able to elope, how long before the facility was aware and acted? When was the resident found and who found the resident? Was the outside temperature or weather a factor? Where the resident was located, were elopement precautions in place? How did the resident exit the facility? What hazards are near the facility? And why did this happen? Was the care plan followed? Staff properly trained to manage exit behaviors? Was the patient supervised? Were there other contributing factors that should be considered? To protect the resident, take immediate action. For example, protect the patient by initiating 1: Convene the team to review the plan of care and assure the appropriate modifications are made to protect from harm. Plan for appropriate placement to protect the resident until locks and keypads are addressed. Address the physical environment issues to secure the unit and keep residents safe until you can address all investigative items contributing to the situation. It is too your advantage to address as many contributing factors as possible while the surveyor is still onsite. Only onsite confirmation of corrective actions justifies a determination to remove Immediate Jeopardy. If corrective actions are instituted before the end of the survey, the IJ is still reported but it includes the surveyor reviewed actions that were taken to ameliorate the risk. Free eNewsletter Subscription Sign up to receive our eNewsletter!

Chapter 9 : Codes, Graphs, and Systems : Richard E. Blahut :

The love affair with tuples and dynamic behavior made the code non trivial and likely is going to cause maintenance issues down the line. It is also quite obvious that this is intended for internal consumption, with very little time or effort spent on "productization".

Bibliographic record and links to related information available from the Library of Congress catalog
Information from electronic data provided by the publisher. May be incomplete or contain other coding.
Introduction to information theory 2. Probability, entropy, and inference 3. More about inference Part I. The source coding theorem 5. Codes for integers Part II. Correlated random variables 9. Communication over a noisy channel The noisy-channel coding theorem Error-correcting codes and real channels Part III. Further Topics in Information Theory: Very good linear codes exist Further exercises on information theory Communication over constrained noiseless channels Information acquisition and evolution Part IV. An example inference task: Exact inference by complete enumeration Maximum likelihood and clustering Useful probability distributions Exact marginalization in trellises Exact marginalization in graphs Monte Carlo methods Efficient Monte Carlo methods Exact Monte Carlo sampling Independent component analysis and latent variable modelling Random inference topics Bayesian inference and sampling theory Part V. Introduction to neural networks The single neuron as a classifier Capacity of a single neuron Learning as inference Supervised learning in multilayer networks Sparse Graph Codes Low-density parity-check codes Convolutional codes and turbo codes Digital fountain codes Part VII. Some mathematics Bibliography Index. Library of Congress subject headings for this publication: