# Understanding cryptographic algorithms through experiential learning

Teaching the inner workings of ciphers such as Advanced Encryption Standard (AES) or Data Encryption Standard (DES) can be improved with the use of visualisation tools. It is proposed that the use of Empirical Modelling techniques to construct models of these algorithms will assist in understanding how they work. This is because the interactive nature of EM allows a cipher model to be used by a student for personal exploration to solidify their understanding and knowledge of the algorithm. Models could also be created for more concrete ciphers, for example a model of an Enigma machine.

The model will be created using JS-EDEN. This is because EM in the browser is easier for people to access since the majority of computers have an internet connection and a browser. As a by-product, this investigation will also analyse any limitations of JS-EDEN when constructing complex models and how it could be improved to overcome these.

## References


W.M. Beynon. *Rethinking programming: modelling heapsort as a case study.* Lecture at UoW Nov. 2012

W.M. Beynon. *Visualisation using Empirical Modelling principles and tools.* AHRC ICT Methods Network Expert Workshop, June 2007

W.M. Beynon and A. Harfield. *Lifelong Learning, Empirical Modelling and the Promises of Constructivism.* Journal of Computers, Volume 2, Issue 3, May 2007, 43-55

Duggins. *Can Empirical Modelling Facilitate Learning?: Modelling Graphs to Assist in the Teaching of AI Search Algorithms.* WEB-EM-4

J. Tao et al., DESvisual: A Visualization Tool for the DES Cipher, *Journal of Computing Science in Colleges,* Vol. 27 (2011), No. 1 (October), pp. 81-89


## Provisional marking weights

40/60 (paper/model)

# Title of Paper

Creating an interactive learning environment powered by a JS-EDEN model to aid in the understanding of vectors.

# Abstract

This paper presents and discusses a JS-EDEN model, codenamed AST that will be submitted to WEB-EM 09 and is designed to aid in the understanding of vectors, an A-Level Mathematics and Further Mathematics topic.

The AST model will consist of both a model of an interactive 2-dimensional vector space and scaffolding (Rogoff, 1990), in this case informative text also displayed in the model, which guides the user's/learner's understanding of the above vector space model and hence vectors. The vector space model will initially be a sparse, simple, microworld (Papert, 1993) and transform to a more complex microworld as the user/learner progress through the scaffolding.

This paper will discuss how Empirical Modelling (EM) techniques and tools can enrich the learning of vectors through interaction and guided discovery, such as how EM can offer a richer web like learning structure of other programs written in an imperative language (Roe, 2003). The paper will also touch upon how these techniques can be applied to other areas of mathematics and indeed, all areas of school education. It will also discuss how AST has gained insight from the methodologies behind sports coaching, as this is often, if not exclusively, an interactive process between the coach and the learner.

# Description

The model submitted to WEB-EM 09, codenamed AST, will be designed to aid in the understanding of vectors, an A-Level Mathematics and Further Mathematics topic.

The AST model will consist of both a model of an interactive 2-dimensional vector space and scaffolding (Rogoff, 1990), in this case informative text also displayed in the model, which guides the user's/learner's understanding of the above vector space model and hence vectors. The vector space model will initially be a sparse, simple, microworld (Papert, 1993) and transform to a more complex microworld as the user/learner progress through the scaffolding.

The areas hoped to be explored the by AST model include:

- What a vector is, using the MEI definition
- Calculating the magnitude of a vector
- Vector addition and subtraction
- Multiplying a vector by a scalar
- The scalar product and orthogonal vectors

The model hopes to separate itself from "traditional" methods of teaching mathematics that emphasizes lectures and note-taking, calculation by hand before using calculators and rote learning (Public Agenda, 2005). A separation from these methods is need because, whilst they are tried and true, the question of their worth has been raised (Public Agenda, 2005). Today's workplace emphasizes flexible and adaptive learners, shouldn't learning methods emphasise this as well (Public Agenda, 2005)?

To mirror this hypothesis, the AST model will adapt and flex to meet the needs of the learner, with learners who pick up the topic easily being allowed to progress through the model quickly and those who prefer to learn by being told are also allowed to progress through the model quickly, reading the scaffolding as they go. Whilst those who enjoy working things out for themselves will be allowed to devote the time to explore the model and come to their own, hopefully correct, conclusions.

It is also hoped that AST can gain some insight from the methodologies behind sports coaching, as this is often, if not exclusively, an interactive process between the coach and the learner.

# Main References:

Alison, C. (2003). Learning through Empirical Modelling.
BCU. (2010). Level 1 Coaching Qualification.
Cronick, A. (2003). Learning through Empirical Modelling.
DCS. (n.d.). Note 3: Vectors. Retrieved 11 15, 2012, from CS131: Mathematics for Computer Scientists II: http://www2.warwick.ac.uk/fac/sci/dcs/teaching/material/cs131/slides_note3.pdf
DCS. (n.d.). Note 7: Basis and Dimension. Retrieved 11 15, 2012, from CS131: Mathematics for Computer Scientists II .
DCS. (n.d.). Notes 5: Linear Combinations and Subspaces. Retrieved 11 15, 2012, from CS131 Mathematics for Computer Scientists II : http://www2.warwick.ac.uk/fac/sci/dcs/teaching/material/cs131/slides_note5.pdf
DCS. (n.d.). Notes 6: Linear Independance. Retrieved 11 15, 2012, from CS131: Mathematics for Computer Scientists II : http://www2.warwick.ac.uk/fac/sci/dcs/teaching/material/cs131/slides_note6.pdf
MEI. (2008). Section 1: Introduction to vectors.
MEI. (2008). Section 1: Introduction to vectors, Notes and Examples.
Papert, S. (1993). The Children's Machine. New York: Basic Books.
Public Agenda. (2005). Teaching Methods: A Public Agenda Citizen Choicework Guide.
Roe, C. (2003). Computers for Learning: An Empirical Modelling perspective.
Rogoff, B. (1990). Apprenticeship in thinking: Cognitive development in social context. New York: University Press.

# Provisional Weighting

The provisional weighting will be 50/50.

# Modelling Aeroplane Boarding Procedures

## Abstract

There are several million commercial flights every year and each one starts off with a few hundred passengers boarding the plane. At the busiest airports planes have strict take off slots that will delay planes either further if missed, which in turn reduces the profit and the reputation of the airline. These factors mean getting passengers onto the plane as quickly as possible is of the utmost importance. Several different procedures have been developed in response to this problem but airlines are unwilling to change from their standard methods that have been in operation for years. These algorithms can be modelled to display each procedure, work out average boarding times and find any potential problems that might arise when using these procedures in practice. The model will look at how these procedures work in a perfect world as well as seeing how different problems that might arise effect the overall boarding time and if they are specific to certain procedures.

## References

http://www.seatguru.com/articles/boarding_procedures.php
http://www.newscientist.com/article/dn20859-test-shows-most-efficient-way-to-board-a-plane.html
http://blogs.discovermagazine.com/discoblog/2011/08/29/the-plane-truth-boarding-by-rows-is-the-worst-possible-way-says-physicist/

## Provisional Weighting

40 Paper
60 Model

# Web-EM 2012

Abstract:

Empirical modeling has strong links with the field of Artificial Intelligence through the shared use of the concept of Agents/Agency. (An agent is something that might change a currently observable situation).
This paper seeks to expand the use of Empirical modeling as used in decision support to encompass automated decision making.
The desirability of this extension is driven by the increasing desire to automate processes and activities contained within the concepts of "Smart" house, work and social spaces etc. Within the "smart" concept there is an expectation that there should be decision making tools which, given users preferences, are capable of negotiating amongst the agents of people within the environment to provide the "least-unpreferable" decision.

In this paper, this will be modeled in terms of a music selection by genre. While this will be expressed in a fairly simple model, the desire is to showcase the decision making concepts possible based upon the empirical model.
Though simple the model does have real world potential application. If all of the people attending a social location, say a public house, expressed their music preference by genre (through their mobile devices, social media etc) the background music or play lists available on the duke box could be determined to provide the music least objectionable to the majority of attending agents. It could also alert others as to locations where their preferred music was being played.

## The Model Design:

The basic design will include three agents and will reference the "avoiding collisions in corridors" model. In its basic form the user will be able to provide input in the form of 3 "genres" (Keywords) and a relevant score for each agent (1,2,3). The system will then determine the least-unpreferable choice and report the score that it has obtained. The least-unpreferable choice is chosen for being the most likely to result in a mutually acceptable choice as people may react more strongly to something they are against than something they are amenable to.

The scoring system will work in the following manner. The scores will go in order of the highest preference; with the highest preference having the lowest value i.e. the lower the score the more desirable it is to the group of users. This method of scoring was selected as being more intuitive to users not used to giving a utility score to choices i.e. "My first preference is Number 1". It is normally the case that a score is given to reflect maximized utility, i.e. the more useful to the user the higher the score; however it can be seen that the scoring could work either way round, depending on the script used by the agent. (In a full deployment consideration could be given to enhancing the model to reflect how the user wishes to score their responses, supported by an appropriate standardisation utility)

The project will explore the difficulties of creating a peer to peer decision system i.e. one in which the agents decide amongst themselves. This may be challenging in a definitive language. There is the possibility of simplifying the task by having a coordinator which would act as an arbitrator to make the decision. This might be achieved through one of the agents being randomly given the role of arbitrator or less eloquently by introducing a third party coordinator and associated infrastructure (someone would need to run the coordinator). Such an infrastructure would also take the model away from the collision avoidance model which was the inspiration.

The model will be built from the modeler's view of the world and their perception of the appropriate results, which will need to be validated. Defining the genres on their own could cause a large number of definitions if the number of genres were increased but the focus is on the decision making process.

References:

~empublic/projects/crowdbehaviourMartin2003
Victor Sanchez-Anguix, Vicente Julian, Vicente Botti, and Ana García-Fornes , 2012, Reaching Unanimous Agreements Within Agent-Based Negotiation Teams With Linear and Monotonic Utility Functions . IEEE Transactions on Systems, man, and cybernetics – Part B Cybernetics, 44(3) 778-792
Beynon, M. 1999. Empirical modelling and the foundations of artificial intelligence. Computation for metaphors, analogy, and agents, 322-365.
Beynon, M and Yung Y.P. 1992 Agent-oriented Modelling for Discrete-Event Systems. Proc IEE Coll. "Discrete-Event Dynamic Systems", Digest #1992/138

# Empirical Modelling Project

## Title
A Model for Interactively Examining Normal Form Two Player Games


## Abstract
The project aims to build a model for interactively examining normal form  games. The model will show the payoff matrix of a two player game, allowing the user to input their own utilities into the matrix. The model will also highlight the dominant strategies for each player, and the Nash equilibria (if any exist) for that game. The user will be able to change the values of the model and see what effect their changes have on the properties of the game. This should act as a learning aid for people who are new to game theory.

I will start the project with very little knowledge of game theory, and use the experience of modelling as a way of learning about the fundamental features of games discussed above. I will structure the report in part as an exploration of using the process of modelling as a method of learning about some subject.


## Preliminary Discussion
I intend to use jseden to build my model, I don't really know how I'm going to implement the algorithm for determining the Nash equilibria yet, as I hope to learn that while building the model. The model will display a table showing the utilities for each strategy chosen by each player, the dominant strategies for each player and the Nash equilibria of the game will be highlighted in the table. The model will have to be constrained to two players so that the table can be displayed in 2D. I also hope to be able to show graphically the relationships between the numbers in the payoff matrix, and how these relationships affect the dominant strategy (as I think that this would be useful to a user).

## Sources to Consult
There are plenty of sources of information on game theory available online. CS407 Agent Based Systems also contains some information. CS409 Algorithmic Game Theory uses 'Algorithmic Game Theory' by Nisan et al. which I've looked at briefly and seems to be a good source.

## Weighting
I will probably opt for the 50/50 weighting.

# A study into potential uses of Empirical Modelling in improving the Efficiency of the Road System

## Abstract

The roads are one of the most important parts of modern day infrastructure. We rely on them to transport goods and people across short and long distances. There are many different types of junctions and roundabouts and the current methods of teaching the Highway Code do not offer the ability to interactively explore and build up a construal of the systems. This study will primarily investigate two things: 1) how we can use the model of a traffic system to aid in driver education; and 2) how a similar model could then be used or built by traffic accident investigators in order to better understand what happened and ascertain blame. The basis for this study will be the existing models “Car Interactions at Roundabouts” and “Modelling traffic flow at traffic light controlled junctions” which will then be augmented with the addition of a model of the car itself, as well as blind spots which occur to the driver.

## Description

The inspiration for this model has been drawn from personal experience where a driver drove into the side of my car while pulling out of a junction. After a complicated claims process involving me drawing several diagrams I eventually won. This made me think that there needed to be a better method of teaching young drivers. Through preliminary research it appears that there are two papers which discuss a similar idea, but neither of the two papers explains blindspots which is where the bulk of this contribution will lie.

The model will be inspired by the two models, noted in the abstract, and written in the “emile” version of JSEden. Where possible the models will be ported to JSEden from the existing models so that a comparison can be drawn between the offline and online tools, and the advantages and limitations of both can be considered. After developing the initial model, it will be augmented through the creation of a model of a car along with the blindspots that are commonly encountered by a driver in the situation in which they will be placed.

The paper shall consider existing methods for teaching learner drivers and evaluate the Empirical Modelling approach employed. It will then draw upon the ideas developed within the model and will be used to compare the numerous aspects of JSEden as well as potential areas for improvement.

# Weighting

It is anticipated that the paper will form quite a large proportion of the assessment and a weighting of 50-60% are expected on the paper.

**Author:** Ivan Leong
**University No:** 0830934
**Course:** MEng Computing Systems

# Finding the optimum strategy in a game of darts

Darts is a game that is normally played in pubs or in professional competitions. To excel in the game one has to be highly accurate in shooting the dart to the highest value sectors of the board (i.e. triple 20). Most players regardless of their skills tend to aim for these targets but it's not always an optimum strategy. This is because a dart throw can be modelled by a 2D Gaussian distribution, meaning that in some cases the player will score the value that he is aiming for and the rest of the time it will land on neighbouring sectors. The board is designed in such a way that high value sectors are placed adjacent to low value sectors (E.g. 20 is adjacent to 1 and 5). In this coursework I plan to develop a model which aims to facilitate players with poorer accuracy in finding the optimum strategy.

Firstly, prior to finding the optimum location of a dart throw, we need to know the accuracy of the player. To determine this, we allow a player to aim n number of shots at the bullseye and we calculate the standard deviation of these shots. The stochastic nature of these shots will be distributed according to the Gaussian distribution function. Hence, we say that a player with a lower standard deviation to be more accurate with shots that land in a smaller cluster around their intended target.

The model I will be developing will represent a dart board where a circular target is divided into 20 segments each with different score values ranging from 1 to 20. There is an 'inner' ring which runs in the middle of the segments and scores double the segment value and an 'outer' edge ring which scores 'triple' values and lastly, at the centre of the board is the bullseye which scores 50 and an outer region which scores 25. Observables can be represented as the different sectors of the board and the dart and dependencies includes the values of these sectors and the total score obtained from throwing the darts. The agent will be the player which throws the dart to the board

Using empirical modelling I am able to use the experiences from a game of darts to create a model in which the outcomes are reflected from those experiences. For example, to model the shooting accuracy of a player, I can have a meter (representing power) with constant changing values (from 0 to 100 and 100 to 0 repeatedly) and is selected when the player releases the dart. While the player is waiting for the right power, the dart can also swerve sideways replicating the motion of unstable hands. Power and instability are two main factors that affect the inaccuracy of a shot.

Once these factors are incorporated in the model, I can then simulate which is the best location for a shot depending on the accuracy of the player. Lastly the model will be able to host a game of darts where a human can play against the computer using the optimum strategies learnt previously.

## References

[1] DataGenetic. A geek plays darts [Online]. (URL http://www.datagenetics.com/blog/january12012/index.html). 2012. (Accessed November 22 2012)

[2] R. J. Tibshirani, A. Price, J. Taylor. A Statistician Plays Darts. JRSS Series A, Vol. 174, No. 1, 213-226, 2011 (Accessed November 22 2012)

## Provisional Weighting

Paper: 40% Model 60%

# CS405 Introduction to Empirical Modelling

## Ninth Warwick Electronic Bulletin on EM (WEB-EM-9)

## "Empirical Modelling as an Experimental Problem Solving Application"

0814909 - MEng Coursework Component

November 20, 2012

**Abstract**

In this project we seek to investigate Empirical Modelling (EM) as a potential medium for accessible problem solving and the study of problem solving methods. The motivation behind examining EM as a medium for problem solving is to draw comparisons with other methods in the fields of Mathematics, Programming and Computer Science. The goal of this study is to create a number of small models that demonstrate the capabilities of EM in educating users in general rules of problem solving. Each of these models will display a different problem which shows a different guideline for solving problems. These chosen guidelines will be a subset of the rules outlined by Problem Solvers and Mathematicians, including Pólya and Mason. We will examine the JS-EDEN environment's benefits as an application with regard to problem construction and flexibility of solutions. We seek to show JS-EDEN as a tool to allow users to experiment and learn from a problem at will and demonstrate that learning in an EM environment can be a more personal way of learning.

# Modelling Study Description

At this point the final state of the models to be made are unknown. Presently we hope to create a different model, of small complexity, to display and educate users about each guideline of problem solving. The list of chosen guidelines is currently in flux but it is likely that the topics of Analogy (Pólya, 1990), Generalization and Specialization (Mason, Burton and Stacey, 2010) will make an appearance. Some of the models may be constructions of examples within Pólya or Mason. We also hope to develop our own original models to display our own thinking and reasoning for new examples. The premise of the models is not graphical fidelity but of ease of exploration. All models will be accompanied with presentation environment slides to suggest routes of exploration, explanation of model features and our own stream of thought. We also hope to use the JS-EDEN specific input features, such as slider bars, to produce a more flexible way of reasoning within a given puzzle or problem than is possible in other environments.

# Primary References

Mason, J., Burton, L. and Stacey, K., 2010. *Mathematical Thinking*. 2nd ed. Harlow: Pearson Education Limited

Pólya, G., 1990. *How To Solve It*. 2nd ed. London: Princeton University Press in Penguin Books

# Provisional Weighting

50:50

**CS405 Empirical Modelling – Abstract and Modelling Study for Assignment 1**

**By**: Judy Palimonka
**Title:** The Operation of Web Tracking in Targeted Advertising
**Date:** 21st of November 2012
**Weighting:** 3 : 7 (paper : model)

# Abstract

Targeted advertising (TA) allows for adapting advertising content to a particular user or an audience of users. It employs a range of web tracking techniques that enable the collection of user data which is then used to adapt the advertising content to the user. Web tracking is fundamental for TA and is accomplished by first, tagging user with a unique ID and then associating user's web usage with this ID. The examples of gathered data can be the history of visited websites, a list of accessed links, click stream, purchases and many more. The techniques of data collection include first-party, third-party and fourth-party tracking. The former technique is utilised by the web publishers who use the collected data in order to advertise their own products and, since they do not share it with others, the process is considered to be relatively safe. Websites employing third-party tracking, on the other hand, allow other websites (trackers) to collect data about each visitor. Fourth-party tracking is usually a result of third-party tracking and takes place when the trackers pass the received data further onto other companies. [3]

The third-party and fourth-party tracking are hidden from the user and there is no effort made from the web publishers' side to obtain user's consent in an explicit manner. The lack of transparency in this area of TA raises many privacy issues among Internet users and privacy enthusiasts. Current research [1, 2] suggests that those aware of web tracking, are concerned about some specific issues such as the amount of collected data, its sensitivity (with respect to users), the stage or time of data collection, the entities collecting data and their intentions. Furthermore, the targeted adverts are not annotated (or vaguely annotated) which makes it more difficult for the user to see the connection between web tracking and advertising. However, this only applies to the people who are aware of the practice. There is still a large percentage of users who are not aware of or are vaguely familiar with the web tracking for advertising purposes.

The lack of visibility in TA is indeed problematic. Firstly, it leaves the public in ignorance and secondly, it does not give them the choice to decide what happens to their data. On the other hand, in order to make an *informed* choice, the public would need to be provided with the practical means of learning about TA. As we could see with the implications of introducing new cookie policy in May 2012 [4], asking user for consent on visit to each website is not an answer to privacy concerns and user's protection. The reason being that it does not guarantee that sufficient information is provided for the user to actually understand tracking and continuous prompts have a negative impact on browsing experience (may even cause users to "ok" the message without giving it any consideration) . A better solution perhaps, would be to engage users with an interactive model that explains TA . In such a way, they could make a decision which would apply to all the websites they visit.

The purpose of this project is to address the above privacy concerns by providing the means to learn about tracking employed by TA. The model will help develop a construal of how web tracking is used in TA and introduce the modeller to the notions of first-party website, tracker, third-party tracking, fourth-party tracking (or sharing), targeted adverts and the relationship between these concepts. Apart from providing the artefacts for the interaction, the model will provide the textual information on how the users can block web tracking. We hope such understanding will enable the modeller to make an informed choice about web-tracking.

# Plan for Modelling Study

**Model:** The Operation of Web Tracking in Targeted Advertising
**Weighting:** 3 : 7 (paper : model)

**Purpose:**

The goal is to let the modeller visualise the flow of data in targeted advertising by demonstrating the operation of third- and "fourth"- party tracking techniques which facilitate targeting users with adverts. The model also helps to understand the impact of web tracking on user privacy (data is passed around without asking for consent).

It would be possible for the model to show how blocking of trackers affects user visibility and the state of targeted adverts while browsing. On the other hand, introducing another concept might make the model more difficult to comprehend therefore this is to decide later.

**Key Observables in Graphical Model**

The following is the interaction console in the model which will comprise the following sections:

- **Upper Area:** A panel displaying the details relevant to the currently viewed web page
  User Visibility: [anonymous, recognised]
  A list of trackers present on website, the attributes of each tracker being:
    Tracker name
    The visited websites where the same tracker was present
    If applicable, the number of trackers that received data from this tracker (fourth-party data sharing)
  The number of targeted adverts

- **Middle Area**: A "browser" showing the address bar and the contents of a "web page"
    Address bar (a form of combo box with predefined addresses)
    Can potentially include targeted adverts with links of other websites

- **Lower Area**: A panel displaying up-to-date browsing progress
    The total number of visited websites
    The total number of trackers

The list of websites will be predefined. In order to access a website, the modeller will need to select it from the list in the address bar. Websites are associated with trackers. Some trackers are sharing user data yet with other trackers which is referred to as fourth-party sharing. The content of each website will resemble real-world page but the URL will be (or at least is meant to be) fictional. There will be a place-holder for an advert on each website. A targeted advert will appear if the following conditions hold:

- User had visited at least one website, say website A, prior to visiting current website, name it: website B
- Both A and B are associated with the same tracker

**Agents**
Website, Tracker, Advert

**Source of Stimuli**
Address Bar, Links in Adverts

**Construal**

The goal of this model is to introduce the fundamental entities in TA and explain how they play role in the process of generating targeted adverts for the user.

The modeller should be able to recognise the following concepts and patterns:

- Targeted advertising is facilitated by web tracking
- The parties involved in web tracking are: website and tracker
- A website passes user data to a tracker or a number of trackers
- A tracker can pass user's data to tracker(s)
- A tracker can be a partner with many websites
- A website can be a partner with many trackers
- A tracker may share user's data with one or more trackers
- User may be recognised by a website even though he/she hasn't been to the website
- (optionally) If tracking is blocked, the number of trackers does not increase and there are no targeted adverts
- Targeted adverts reflect on previously visited websites

**Similar Applications and Related Research**

The operation of this model resembles the operation of “Detective” - the program I developed as a part of my final year BSc project. Both, the program [5] and user manual [6] can be accessed from the web for reference.
The applications such as “Ghostery” [7] or “AVG do not track” [8] are similar to Detective but only to some extent as they do not recognise fourth-party sharing, neither annotate targeted adverts and also they use different methods for recognising trackers. Also all these programs use browsing in real-time to show how trackers. Contrary to these programs, the model will allow for accessing a certain number websites with predefined content.

**Draft of Graphical Interacting Model**

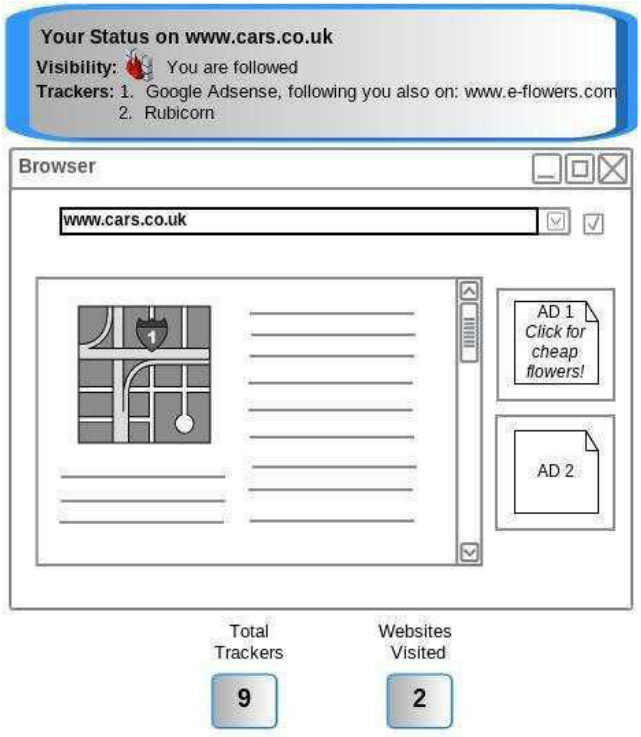


This is only a rough sketch of how the model could look like. The content will consist of text relevant to the URL, e.g. cars for sale available at URL contained in the model below. Targeted adverts will be linked to user's previous activity. In Trackers fields, the model is telling user that Google Adsense has been tracking user on www.e-flowers.com. To show how targeting is related to browsing activity, AD 1 is showing an offer related to flowers.

## References

[1] OFC, *A market study: Online Targeting of Advertising* , London, UK: OFC, 2010.
[2] IAB, *A guide to Online Behavioural Advertising and Prices* , London, UK: IAB, 2011
[3] J. Mayer, J. Mitchell, *Third-Party Web Tracking: Policy and Technology*, Stanford, US: Stanford University, 2012
[4] ICO, *New EU cookie law (e-Privacy Directive)*, available at
http://www.ico.gov.uk/for_organisations/privacy_and_electronic_communications/the_guide/cookies.aspx
[5] Palimonka J., *"Detective – Increasing Transparency of TA"*, Mozilla Firefox, 2012 available at
https://addons.mozilla.org/en-US/firefox/addon/detective-increasing-visibi
[6] (not published, provided here only for reference) Palimonka J., Project Report, 2012 available at
https://docs.google.com/open?id=0BwMaPNkfLHPYaDhrekZSTWNoSGM
[7] Evidon, *"Ghostery"* application, available at http://www.ghostery.com/about
[8] AVG, "*Do not track*" application, available at http://www.avg.com/gb-en/do-not-track

# Empirical Modeling Project Proposal

Stefania Papaconstantinou

## Proposed Title

Migration of the *Hunt The Wumpus* game model into JS-EDEN.

## Abstract

This project will aim to migrate the existing model of the *Hunt The Wumpus* game implemented in EDEN by Graeme Cole, to JS-EDEN. JS-EDEN is a definitive browser environment that supports the %eden notation, but does not include other notations such as %donald, %scout and %aop that the original model uses. JS-EDEN, on the other hand, enables the user to implement and use the model in a familiar object-oriented way. Implementing the game in JS-EDEN will therefore provide a good example of a model that is more suitably modeled using the principles of object orientation. The *Hunt The Wumpus* game is based on an agent that consists of possible actions that it can perform, and a series of observables which it has some agency over. The project will explore the nature of the observables and dependencies in the current EDEN model with respect to the literature of communicating construals and aim to translate them into the equivalent JS-EDEN artifact. The *Hunt the Wumpus* game is a classic example of some simple Artificial Intelligence concepts, as it is based on making inferences about knowledge gained from perceptions, and it would therefore be a useful learning model.

## Preliminary Description

The *Hunt The Wumpus* game is a grid-based game that was originally conceived and written in BASIC by Gregory Yob. The game's grid has the shape of a kite-shaped dodecahedron, which represents a cave, with the edges representing the connective paths between the nodes, which are rooms within the cave. The aim of the game is to kill the "Wumpus", a beast that eats the agent if it moves to its room. The agent has to achieve that without getting eaten by the Wumpus, before running out of arrows and at the same time, whilst avoiding other hazards, which include falling into a pit and being moved by bats. The agent has no prior knowledge of the contents of each room but can sense the Wumpus and other hazards, if they are near.

Prior to starting my implementation of the model in JS-EDEN, I will first study how the existing model is built, as it will be important to include all the necessary observables and dependencies that can be found in the original model. Moreover, as the model also aims to achieve a learning goal, as it requires the player to act based on knowledge gained from perceptions, I will be looking to identify how the model addresses this goal and to what extend it is achieved. This study should hopefully not only provide enough information for the remodeling of the game, but should also show any shortcomings of the original model, which I will aim to resolve.

After the implementation, I will be evaluating my model compared to the original, but also to other implementations of the *Hunt The Wumpus* game that exist around the Web. I will also aim to test with actual players and make changes depending on how successful the model was in demonstrating this AI concept.

**References**


1) G. Cole, Hunt the Wumpus: an Empirical Approach, Undergraduate fourth year module, 2005, Report URL: <http://empublic.dcs.warwick.ac.uk/projects/wumpusCole2005/wumpus.pdf>, Project URL: <http://empublic.dcs.warwick.ac.uk/projects/wumpusCole2005/>
2) J. Pavelin, Wumpus World, Undergraduate third year project, 2002, URL: <http://empublic.dcs.warwick.ac.uk/projects/wumpusPavelin2002/>
3) T. Monks, A Definitive System For The Browser, 2011
4) S. J. Russel and P. Norvig, Artificial Intelligence: A Modern Approach (3rd edition), Pearson, 2010
5) G. Yob, Hunt The Wumpus: The Genesis of Wumpus, 1976, URL: <http://www.atariarchives.org/bcc1/showpage.php?page=247>


**Proposed Assessment Weighting:** 50%:50%

# Empirical Modelling Proposal

Adam Riley
Deparment of Computer Science
University of Warwick

## CONTENTS



## 1 ABSTRACT

I propose to use the EM modelling environment to define a state based Artificial Intelligence capable of simulating human emotion. By defining each state as a pre-determined emotion, the model will, using information provided by the user, infer qualities about its environment, and dynamically adjust its current 'temperament'.

As some of the environmental factors may directly affect the simulated being, such a temperature, time of day or thirst, there will be many interdependencies between not only the internal emotional states of the model, but also the observables displayed to the users.

The area of human emotional simulation is often confined to replicating observable measurments from a human subject, in this instance, I will construct a model, using pre-defined emotional states, which will act as an independent agent.

By using this model to determine the bahaviour of a single agent under any given enironmental conditions, it will then be possible to determine appropriate interactions between two agents, undergoing any combination of emotional states.

## 2 MODEL OVERVIEW

This model will need to accurately represent the complex notions of an artificial intelligence in a suitable manner, such that any user can easily determine the current state of the agent at any gven time.

In order to ease the viewing process by any users, each state will be colour-coded. Each of the variable environmental factors available to the user will be displayed through a GUI with features such as numerical displays, sliders and execute buttons. It may also be necessary to visually represent the environment, with pictures or colours.

I intend to approach this problem with a focus on dynamic states. As such, each state in the system will always be active and attempting to control the agent at once. By representing the level of control each emotion has over the agent at any given time as a progress bar, the dynamic nature of the emotions will be exposed to the user.

It is widely considered that there are at least six basic emotions: "anger, disgust, fear, happiness, sadness and surprise." [1]

It is important thateach of these emotions, and any ancilliary emotions included in the model have adequate definitions, and can affect each others control, in order to keep the system fluid. This will be accomplished by representing each emotions overriding effect on others as a dependency network, with stronger emotions having more of an effect when altering weaker ones.

Possible extensions to this model could be to introduce the notion of multiple agents. As each agent would be calibrated differently, to reflect the asymmetric nature of personality, each agent could be in a different state to the others in a given enironment. It may be of interest to then model their effect on each other, such as a single agent in a 'happy' state, interacting with multiple agents in the 'sad' state.

## 3 Proposesed Weighting

60% : Model
40% : Documentation

## References

[1] Antje Herbon Christian Peter. Emotion representation and physiology assignments in digital systems. *Interacting with Computers*, 18-2:139–170, 2006.

[2] Pasqualino Ferrentino Shuji Hashimoto Kenji Suzuki, Antonio Camurri. Intelligent agent system for human-robot interaction through artificial emotion. *IEEE International Conference on Systems, Man, and Cybernetics*, 2:1055–1060, 1998.

[3] Picard. R.W. Synthetic emotion. *IEEE Computer Graphics and Applications*, 20-1:52–53, 2000.

[4] Juan D. Velsquez. Modeling emotions and other motivations in synthetic agents. Master's thesis, MIT Artificial Intelligence Laboratory, 1997.

# CS405 - Modelling Study Proposal: Wireless Sensor Network

Daniel Robertson (0910210)

27th November 2012

## 1 Abstract

The recent emergence of Wireless Sensor Networks (WSNs) - due to both their decrease in cost and size, and increase in battery life - has seen their rising application in many real world problems. These applications fall into 3 categories: Military Applications - information collection such as enemy tracking, battlefield surveillance or target classification are all areas being researched actively, Environmental Monitoring - both indoor monitoring for emergency services and outdoor monitoring with applications to ecology systems and agricultural processes and Logistics - ranging applications from asset tracking using RFIDs to monitoring a company's industrial process and heavy machinery. As can be clearly seen the application domain for WSNs is vast, and as such it has become a very active research area.

However, when working with WSNs there are several considerations that a developer must make. In addition to the complexity added by making a system wireless, WSNs feature a different set of constraints, such as the requirement to minimise energy consumptions and a limited wireless range. These present new challenges when designing algorithms and software for these devices. In an attempt to solve these problems, bespoke software will need to be developed to operate the WSN hardware. As sensor networks are a distributed system, this increases the complexity even further, as the software can suffer from dead-lock, live-lock and other concurrency issues. My aim is to create an empirical model of a WSN to assist in understanding how all these factors come together.

## 2 Preliminary Description

The subject of my model will be a Wireless Sensor Network featuring 10-20 nodes (also called motes); this is considered a very small network, as "real life" WSNs can feature hundreds or even thousands of motes. However, this number of motes should prove sufficient to demonstrate the key principles of WSNs while not being large enough to be overwhelming and thus impede understanding. The model will represent several features of a network's setup and operation, including:

- **Clustering** - the setup process in which each mote chooses a parent (or clusterhead) to which its messages will be sent; this parent then aggregates the messages before sending them on to reduce network traffic.
- **Neighbour detection** - this has many varieties, but the most common (and the I will model) is establishing a list of all nodes within a given nodes one-hop neighbourhood.
- **Data generation/transmission** - once the network is set up, motes are typically used to gather data using their sensors and pass this back to some base station (sink) node. In some applications, any node in the WSN may send a request for data, either to one specific node or to (some subset of) the entire network. Again, there are many ways to implement these, both with and without a clustered network. Minimising the number of messages sent (thereby extending battery life) is one of the key research areas regarding WSNs.

- **Message logging** - in case of error, it is sometimes convenient to store the most recent messages, which can be later analysed to better understand the problem.

To facilitate these, the radio range of the motes will need to be considered; this typically refers to the "logical" range rather than the physical distance from the node, as real-life deployments often feature some conditions (obstacles, etc.) which limit the actual range of the motes' radios. Of course, in a simulation or model of a WSN, the logical and physical ranges can be considered equivalent.

# 3 Resources

As this topic is related to my fourth-year project, I will be basing parts of this model on my first-hand experiences of developing a WSN. However, it is also an active research area with many academic papers relating to all manner of sub-areas of the WSN field, some of which may prove useful in understanding the concepts of WSNs and therefore aid in the modelling process. There are two primary choices in operating system for a WSN; TinyOS and Contiki. As my fourth-year project uses Contiki, I will also be referring to this OS (insofar as generalisation of concepts can't be avoided) in my model; specifically the Cooja WSN simulation environment which is only compatible with Contiki, as it will provides key assistance in how to visualise a WSN. I also have limited access to a set of actual wireless sensor motes, experiences with which will inform how a purely theoretical model could be adjusted to include issues imposed by reality.

# 4 Weighting

I tentatively anticipate the weighting of my submitted paper against the finished model to be a 50:50 split.

# SVGJSEden: extending the observable capacity of a drawable

Bernard Sexton

21st November, 2012

**Abstract**

jsEDEN is a definitive environment for the browser modelled on the principles of observables, dependency and agency of Empirical Modelling, supported by the EDEN (Engine for Definitive Notations) environment. A significant benefit of jsEDEN is it's compatibility with a number of browsers and contexts. However, the expressiveness of the models it defines are limited by the range of drawables and their parameters in the browser environment. This proposal for WEB-EM9 will explore the generalised modelling practice already employed by jsEDEN and how it already successfully supports the notion of 'state-as-experienced' [7] through allowing for gradual refinement of models through observation and action, allowing one's own understanding to develop from their interactions with a given artefact.

The project will then consider an SVG implementation of the jsEDEN environment in pursuit of enhancing the notion of "communicable representation" or "construal" conveyed by models defined in the browser environment. The motivation for doing this is grounded in improving how models are construed by the user through improvements of artefact representation in jsEDEN. SVG is a scalable vector graphics format defined using conventional XML, allowing elements drawn using the notation to be transformed and manipulated via Javascript. The extensible nature of SVG allows greater scope for defining the agency that effect the observables and dependencies that ultimately determine how a model is to be construed. The identification of differences between the current HTML5 implementation and an SVG variant will be highlighted by modelling an existing jsEDEN project to identify how artefact representation affects the overall construal of a model.

# 1 Preliminary Description

The power of a visualisation allows the user to directly construe how changes to the observables and dependencies affect a given environment or model. For a model or environment to be construed in a way in which the author intended, it requires effective re-evaluation of it's state in order to maintain the original referent while supporting the construals that the model is communicating. jsEden already supports a basic set of drawables for model representation but prescribes attributes that are used to describe those drawables. Arguably, this limits the capacity of what can be construed through the interaction process with models that use these particular drawables. The project will research into the current state of the *Scalable Vector Graphics* standard alongside other web technologies being used for drawing in the web browser. These technologies are subject to change, so it is important to distinguish how enhancement in artefact representation is an important long term goal for construal comprehension in jsEDEN. The fundamental differences that exist between the single *eden* notation implemented in jsEden and the mulitple notations (such as donald, scout, sasami etc) implemented in tools such as tkEden will be examined.

For a definitive browser environment to communicate the representation anticipated by the author and the modelled environments, the tools used should be capable of expressing a number of modelling contexts adequately. The current use of HTML5 canvas as a platform for visualising these contexts is limited by poor browser support, a narrow range of properties for the described objects and inefficiencies in dealing with changes to the properties of the canvas due to the requirement to redraw after every update invoked by the jsEden interpreter. The benefit of SVG lies in the freedom of expression in the attributes used to describe SVG elements as well as it's exposure to the Javascript document object model meaning the state of SVG elements can be freely altered without re-evaluation of the entire element or the environment.

The project will consider how this is currently achieved in existing tools and aim to implement and improve, where possible, particular drawable components to give better representations of new and existing models in SVG. SVG in the jsEden context will form part of an existing HTML5 environment, therefore the project will consider how these web standards can be combined to provide a more robust platform for modelling and a better mechanism for construing observables, dependency and agency. A discussion of SVG as a mechanism for supporting construing of models will be considered alongside simple models recreated using the SVG based drawable library. The successful construal of a particular model is particularly important in the context of learning so models in this domain, such as Jess Nixon's *Eightball* [4] will be revised using the SVG based solution to see whether improvements in representation allows these models to be construed in a way which better represents the reality they propose.

## 2 Works to be Referenced

[1] Monks, T. (2011). *"A definitive system for the browser"*. Unpublished master's thesis. University of Warwick, Coventry, UK.

[2] Beynon, M; Myers, R; Harfield, A. (2009) *Web Eden: support for computing as construction?* In: International Conference on Computing Education Research (Koli Calling), Koli National Park, Finland.

[3] Beynon, M (2011) *Modelling with Experience: Construal and Construction for Software* In Ways of Thinking, Ways of Seeing Automation, Collaboration, & E-Services Volume 1, 2012, pp 197-228.

[4] Nixon, J (2011) *2-Ball & 7-Ball Eightball* - jsEden Model.

[5] W3C. (16 August, 2011). *Scalable Vector Graphics (SVG) 1.1 (Second Edition) Draft Specification* [Online] (url http://www.w3.org/TR/SVG/). (Accessed on 11/11/2012).

[6] Ward, A. (May, 2004). *Interaction with Meaningful State: Implementing Dependency on Digital Computers* Unpublished PHD thesis. University of Warwick, Coventry, UK.

[7] King, K. (May, 2004). *Uncovering Empirical Modelling* Unpublished Masters thesis. University of Warwick, Coventry, UK.

Other works may be referenced as the project progresses.

## 3 Assessment Weighting

Paper 50%, Model 50%

# Modelling the positioning of players in a five-a-side Football game

Amit Shah

## 1 Abstract

Football is traditionally played with teams of eleven players, however a much more common version of the game involves only five players on each team. These games are played usually at non-professional levels, and as a result, there is little exposure to different formations and tactics, instead relying on scaled down tactics of full teams. It is often seen that players will attempt to build a team based on their own abilities, rather than working together to optimise their layout on the pitch.

In order to maximise the efficiency of the teams, players must continually reposition themselves based on the gathered information from the field. For instance if a player at the back of the pitch has the ball, the remaining players should fall back, and provide support to that player, to be able to receive passes of the ball.

This project will attempt to model a team of five players, in various formations, and then allow the rearrangement of these players. Rearranging a player should also cause the remaining players to adapt their positioning, dependant on the moving player's position.

## 2 Model Description

This project will involve building a constual that demonstrates the positioning of a five-a-side football team on a pitch. These players can be placed in various formations, and the model will provide options for some 'standard' formations via the use of buttons. Displaying the roles of these players is important to a user trying to use the tool to develop their own method of play, this can be done using simple letters alongside the object representing the player. The letters refer to the player's main duties within the game, such as striker or defender. The constual itself will allow the manipulation of players into positions outside of the standard formation, such as placing a player further north in the pitch. In response to this rearrangement, the players that the moved player depends on will be rearranged to show the positions needed to provide support to the original player.

Given certain time limitations, it would have been interesting to be able to produce a constual that shows how players should move, depending on both the player with the ball, and players on the opposing team. However I feel that modelling how one team should react should be enough for a 'real life' situation, as traditionally in other sports (such as American football), each play is defined by how the current team should move around the field.

The observables in the model will be the players themselves. There should not be a need to model the goalkeepers positioning, as they are usually static throughout the game. It may also be useful to

consider the ball as an agent, as it's positioning (when not with a player) should also determine how players should be moving through the pitch (it would be counter intuitive for players to be following another, going in the opposite direction to the ball).

The dependencies will model how the interact with each other, and allow the visualisation of team tactics in the game.

The main users of this model could be people trying to improve their own five-a-side football teams, by introducing new tatics. It will help them learn how they should be positioning themselves, and improve their overall game.

# 3 References

## Previous Models

http://empublic.dcs.warwick.ac.uk/projects/footballTurner2000/

## Tools

http://www.dcs.warwick.ac.uk/ empublic/js-eden/eden.html

David Walton November 2012

# Empirical Modelling and Guitar Notation

**Abstract & Outline of Modelling Study:**
There are several different musical notations available to guitar players, and each has its own characteristics. A piece of written music could be considered as a set of observables, and dependencies could be used to describe the connection between different notations. Notation is also intimately connected with agency, as different forms of notation describe components of music at different levels of detail. This means that, if a different notation is chosen, the freedom the player has to interpret the piece, and hence the player's agency, subtly changes.

The types of notation considered in this paper will be standard musical notation, fretboard diagrams, tablature and chord lists (see below). The connection between them will be empirically investigated via the creation of a model in JSEden. This model should include visual representations of each of the different forms of notation, and allow the user to view a number of examples of chords and short pieces of music in each of the different notations. It should also be possible to play the written music, and at the same time as parts of the piece are played, corresponding elements of notation should be highlighted.

In the accompanying paper, the connections between these forms of notation will be discussed in more detail, as well as their connection with the act of guitar playing. All forms of notation are different ways of converting the writer's informal, hermeneutic interpretation of a piece of music into a formal artefact. Other guitarists can then use this to develop their own personal interpretation, which they can then communicate informally to others through playing. The degree to which the player's interpretation differs from that of the writer depends in part on the notation used, and this connection will be explored in more detail in the paper.

There are a number of potential difficulties in developing dependencies to connect different forms of notation, as each preserves different aspects of the act of playing a piece of music. For example, standard musical notation fully describes the rhythm and melody of a piece of music. However, the same note may be present at a number of different positions on a guitar's fretboard, meaning that conversion to tablature is not well-defined. Conversely, guitar tablature contains very limited rhythm information, meaning that conversion from tablature to standard notation is similarly ambiguous. The created model will aid in exploring the connection between these notations further.

**Examples of Notation**
The image below is a D major arpeggio, written in standard notation above, and guitar tablature below. The standard notation describes the pitch of each note, and the rhythm in which they are played. Reading from left to right, the tablature illustrates the fret numbers of each note played in chronological order, and the horizontal lines correspond to the strings on a guitar. Written as a chord list, this would simply be “D”. To the right is a fretboard diagram. Note that finger numbers are also included. Note that fretboard diagrams and chord lists do not describe the action of the player's plucking hand (generally the right hand).

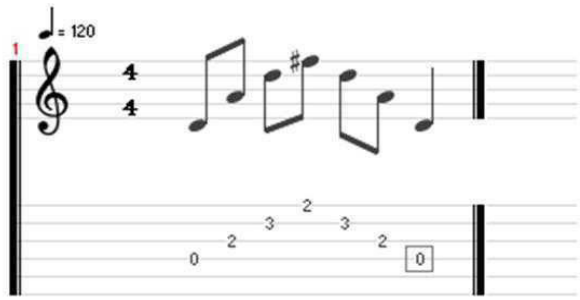


David Walton November 2012

**Preliminary Weighting: 50%:50%**

**List of References:**

The "Music Program" model by David Wai (2000):
http://empublic.dcs.warwick.ac.uk/projects/musicWai2000/

**A number of papers in the EM archive relating to music and EM:**

From formalism to experience: a Jamesian perspective on music, computing and consciousness, Beynon 2011
Mathematics and Music - Models and Morals, Beynon 2006
Dissolving dualities in mind, music and mechanism, Beynon 2006

**A number of papers related to the guitar and notation:**

Laurson, Mikael, et al. "From expressive notation to model-based sound synthesis: a case study of the acoustic guitar." Proceedings of the 1999 International Computer Music Conference. 1999.
*Outlines a system for modelling acoustic guitar sound. Part of this involves extending standard notation to include expression, vibrato and fine timing information.*

Radisavljevic, Aleksander, and Peter Driessen. "Path difference learning for guitar fingering problem." Proceedings of the International Computer Music Conference. Vol. 28. sn, 2004.
*Outlines a model for resolving the ambiguity in interpreting fingering from musical notation.*

Tuohy, D., and W. D. Potter. "A genetic algorithm for the automatic generation of playable guitar tablature." Proceedings of the International Computer Music Conference. sn, 2005.
*This paper outlines a genetic algorithm to automatically generate tablature from musical notation.*

Tuohy, Daniel, and W. Potter. "Generating guitar tablature with LHF notation via DGA and ANN." Advances in Applied Artificial Intelligence (2006): 244-253.
*Another method for tablature generation from musical notation.*

Macrae, Robert, and Simon Dixon. "A guitar tablature score follower." Multimedia and Expo (ICME), 2010 IEEE International Conference on. IEEE, 2010.
*A system for interpreting guitar tablature, and playing back the result whist following the tablature.*

**The following papers are related to the act of playing the guitar:**

Motokawa, Yoichi, and Hideo Saito. "Support system for guitar playing using augmented reality display." Mixed and Augmented Reality, 2006. ISMAR 2006. IEEE/ACM International Symposium on. IEEE, 2006.
*A system for guitar teaching using AR, including overlaying a virtual hand model to demonstrate left hand positioning.*

Heijink, Hank, and Ruud GJ Meulenbroek. "On the complexity of classical guitar playing: functional adaptations to task constraints." Journal of motor behavior 34.4 (2002): 339-351.
*An analysis of the behaviour of classical guitar players.*

# Empirical Modelling of Wireless Sensor Networks

Joseph Yarnall

November 20, 2012

**Abstract**

Wireless Sensor Networks (WSNs) are a relatively new area of research and one which provides interesting and complex challenges. A WSN consists of any number of small battery powered Motes which usually contain sensors such as; temperature, humidity and light. The Motes are also built with wireless transmitters and receivers but because of their energy constraints transmission range is limited; therefore, the Motes in a WSN must work together to relay messages back to the user which is usually referred to as the sink. The goal of a WSN is to remain operational for as long as possible all the while completing the distributed task given to them such as evaluating predicates or relaying values to the sink; this means factors such as mote positioning can have a dramatic affect on the networks effectiveness.

Due to the complexity of calculating the energy efficiency in a WSN, and how even slight changes in network layout can have a massive affect on the network, I have decided to try to model this scenario for my project.

This report will explore the principles, techniques and tools taught in the Introduction to Empirical Modelling module at Warwick University and consider how such an approach can be used in conjunction with Wireless Sensor Networks to construct a model of the dependencies, agencies and observables of a such a system.

# 1 Model Description

My project aims to model the operation of a wireless sensor network deployed in an environment with the task of evaluating a predicate; if the predicate is broken then, the node will report the result to the sink. The network will consist of 9 Motes and a sink Node, the user can arrange the Motes in any layout they wish and can set the predicate to evaluate the statement: TEMP $\geq 40C$ where the temperature is variable by the user.

Once the simulation is started the user will see as the Motes pass messages between one and other. The user will be able to alter the temperature at each node to induce messages in different parts of the network. Then the user will be able to see how this affects the battery life of each Mote in real time. Finally, the simulation will end when all the nodes run out of power and the user will be able to reset the simulation back to the beginning and try a different configuration of Motes.

While the simulation is running the user will be able to alter the configuration of the network to enable them to explore the interactions and dependencies of the model.
There are several dependencies at work within my model, I will now describe the key ones that I aim to model. Firstly, the dependency between the predicate to evaluate and the temperature at a node. Secondly, the dependencies between message sends and receives and battery power. Thirdly, the distance between nodes determines if nodes can communicate with each other. Finally, the dependency between the predicate to be evaluated and sending a message.

There are several ways in which I would improve my model if there was time. Firstly, I'd allow for more complex WSN algorithms to be run on the network, for example the current algorithm is very dumb and makes no effort to create clusters in order to minimise on message sends. I could create a selection box at the beginning of the simulation where the user could select the algorithm they wish to run as well as network configuration.

Modelling Environment: JS Eden

Marking Distribution: 60% Model, 40% Report

# References

[1] JsEden

[2] Tim Monks. A Definitive System for the Browser.

[3] Al-Karaki, J.N. and Kamal, A.E. Routing techniques in wireless sensor networks: a survey Wireless Communications, IEEE.

[4] Chandy, K.M. and Lamport, L. Distributed snapshots: determining global states of distributed systems, ACM Transactions on Computer Systems(TOCS).