Optimisation of BFWA Networks using Emergent Intelligence

A.A.Wade[1, 2], S.Hurley[2], S.M. Allen[2], R.K. Taplin[2], K.H. Craig[1]
[1]Radio Communications Research Unit,
Rutherford Appleton Laboratory
[2]Centre for Intelligent Network Design,
Dept. of Computer Science, Cardiff University
a.a.wade@rl.ac.uk

Key Words
BFWA, Wireless Network Optimisation, Automatic Planning Tool, Meta-Heuristic Optimisation, Emergent Intelligence Systems.

Abstract

This paper presents a detailed account of a novel scheme for finding profit optimal BFWA networks, extending an existing network optimisation tool, and using the principles of emergent, self-organising systems. We describe how populations of agents representing potential users and base sites will disseminate and react to 'local' information to optimise global objectives. The use of two distinct types of agent entity allows the multi-objective profit/coverage nature of the optimisation to be satisfied. Preliminary results are presented to indicate the potential of the scheme.

Introduction

Broadband Fixed Wireless Access (BFWA) networks are becoming an attractive alternative to cable in offering low cost high speed data services, telephony and video to residential and business users. To compete successfully with the available alternative technologies, the effective planning and design of the efficient networks is crucial.

ECHO and the Network Model

To achieve a successful design for a BFWA network, a planning tool must be able to:

- Select base station sites from a number of candidate locations.
- Select and configure the equipment to be installed at these sites (including the frequency assignment plan), to achieve certain coverage or service criteria.
- Select and configure the equipment to be installed at the subscriber (user) premises.

ECHO is one such network planning tool, (developed at Cardiff University for the European Union Information Societies Technologies (IST) project EMBRACE - Efficient Millimetre Broadband Radio Access for Convergence and Evolution [1]), which attempts to optimise and automate all aspects of this process. The tool uses a Tabu-Search[2] type meta-heuristic to produce a profit-optimised network plan. ECHO uses path loss information generated by the Rutherford Appleton Laboratory's Rapid Pipeline Development (RPD) 3-D ray tracing software tool. More information about the network model and the propagation model used in RPD and ECHO can be found in [3],[4],[5].

Emergent Systems

In nature colonies of Harvester ants constantly adjust the number of workers actively foraging for food. No individual ant has the cognitive capacity to assess the colony's food needs, and each has only a meagre chemical vocabulary. However, as a colony, these ants can find the shortest distance to a food source and then prioritise that source, depending on distance and ease of access.

This is an example of swarm logic or *emergent* behaviour - the ability of complex systems to emerge from numerous simple components. The organisation of the colony is carried out from the bottom up, as opposed to the top down organisation our society is more accustomed to. Global information and individual actions determines local actions. In ant colonies and all emergent or self-organising systems, local information and collective action produces global behaviour.

This behaviour has been harnessed in a number of combinatorial optimisation schemes [6],[7],[8], although these approaches are somewhat different to the one detailed in the next section.

An Emergent Intelligence Scheme

The scheme detailed in this paper attempts to recreate this emergent intelligence, through the use of two tiers of semi-intelligent agents. Each agent within the system is based on an actual part of the network, namely the potential *users* of the network and the potential *base-sites*. Every agent will obey simple rules governed by an analysis of local information that they are privy to. By having two distinct, though co-operating, breeds of agent, with opposing goals, the optimisation can drive towards dual objectives simultaneously.

The optimisation proceeds in an iterative manner, with three phases per iteration:

- User-Phase
- Site-Phase
- Control-Phase

The User and Site Phases are best described in the context of their agents. With reference to the three planning needs described previously, it is the User-Agents that select the base-sites to be used, as a consequence of their actions in the User-Phase, and during the Site-Phase, the Site-Agents choose and configure the network infrastructure. During the Control-Phase the planning tool stores the best network found so far as the optimisation proceeds.

User-Agents

User-Agents represent the potential users of the network and so are concerned with providing a service to their user. They are defined as follows:

U_A is defined as the Set of User-Agents.

$u_{Ai} \in U_A$ is an individual User-Agent defined as follows:

$u_{Ai} = \{u_i, d, r, s, c, Nu_i\}$ where:

$u_i \in U$ represents the user parameters from the network structure[3],

$d \in [0,100]$ denotes the User-Agents Distraction level,

$r \geq 0$ denotes the User-Agents Resentment level,

$s \in \{0,1,2,3\}$ denotes the User-Agents current State,

$c = [0,2]$ denotes the relative charge ratio,

$Nu_A = \{u_{Aj}, \ldots, u_{Ak}\}$ denotes the set of Neighbours

During the User-Phase the main concern of each User-Agent is the petitioning of Site-Agents. User-Agents petition the Site-Agents in order to obtain service from them. Each User-Agent then will be in one of four States: *Petitioning*, *Covered and Petitioning*, *Served* or *Dormant*.

One important aspect of swarm logic and emergent systems is the notion of reaction to local information. Each User-Agent has a set of neighbours with whom they can communicate. This set, Nu_A, is defined as all the User-Agents that are positioned within the area bounded by a circle of some radius, R, centred on user u_{Ai}. Each iteration, every User-Agent makes an assessment of its current state and go on to inform their neighbours :

- I am petitioning site *x* for service.
- I am covered by site *x* and petitioning for service.
- I am served by site *x*.
- I am dormant.

Each User-Agent makes a survey of the information received and combined with personal information, three parameters are calculated every iteration – *Distraction*, *Resentment* and *Charge Ratio*.

Each User-Agent keeps track of which Site-Agents their neighbours are talking to. The percentage of a User-Agent's neighbours that are not talking to that User-Agent's preferred Site-Agent is called the level of Distraction, *d*. A User-Agent can be said to be highly distracted if a large proportion of its neighbours are directing their attention elsewhere.

User-Agents respond to the activities of their neighbours, but also to the activities or lack thereof of their chosen sites. A User-Agent is justifiably unhappy with a Site-Agent that is paying it no attention, and if this continues over time then the level of unhappiness with the current solution (from their selfish point of view) grows to significant levels. This unhappiness is called Resentment, *r*, and is measured in the following way:

- In State 1(petitioning) r= t,
- In State2 (covered) r= t/10,
- In State3 (served) r=t/100,

where *t* is the time spent in the current state

One of the problems found early in the implementation of this scheme was the willingness of User-Agents to 'prop up' an inefficient sector. In order to assist the scheme in avoiding this, the User-Agents were given a way of assessing the Base-Site they are receiving service from. This is achieved through the use of a relative subscription charge ratio, c. As in the standard ECHO model each user is offering a certain amount of revenue, (network revenue symbol), to the network in exchange for service. In the AgentOpt scheme this is extended to also include a subscription charge that reflects the number of users served by the particular Site-

Agent. The Site-Agent asses its total costs and divides this number by the total number of User-Agents served, to give an income per User-Agent necessary for the Site-Agent to be profit neutral. The User Agent receives this charge and compares it with the agreed revenue charge.

The Charge Ratio, c, is then defined as the fractional percentage of the charge compared to the agreed subscription rate, limited to give a value between 0 and 2. It is worth being clear that this value is only used by the User-Agent in order to determine its next action. The Site-Agents and indeed the global network income levels considered during the Control-Phase only 'see' the agreed subscription rate.

The User-Agent then takes this percentage into consideration when deciding to change state. This can assist in the migration of users between sectors as explored in the Site-Agent section later.

The probability of a Site-Agent change, P(*sc*), is determined by the User-Agent's level of displeasure and the amount of neighbours behaving differently so that,

$$P(sc) = r*d*c$$

A uniformly distributed random number, *n*, such that 0 <= *n* <100, can then be generated and compared with P(*sc*) so that:

If *n* < P(*sc*) the User-Agent determines to change it's preferred Site-Agent.

The new Site-Agent is chosen from the two most popular in the User-Agent's neighbourhood. If the current preference is the most popular, another random decision, weighted by the percentage popularity of the Site-Agent in question, is taken to determine whether this Site-Agent is worth persisting with.

As can be seen then, the User-Agents levels of Distraction and Resentment drive the optimisation into investigating new solutions. These are the feedback loops mentioned previously. So, User-Agents in Site-Agent dense areas, with high levels of distraction, are more unstable. If a User-Agent has no real choice about the Site-Agent then *d* is low – perhaps zero – and there is no likelihood of a change. The same applies to the resentment levels. User-Agents only have high levels of resentment when ignored for a long time and receiving no answer from the Site-Agent in question, then a change is deemed necessary. Resentment does grow over time even in the User-Agents that are served. This enables the optimisation to try different solutions after a long enough period that the current sector configuration has a chance to get as good as it can get.

Site-Agents

Site-Agents represent the potential base station sites in the network and are concerned with being profitable. Global profit optimisation comes from the competition amongst the individuals within the system.

The Site-Agents are defined as follows:

BS_A is the set of Site-Agents,

$bs_{Ai} \in BS_A$ is an individual Site-Agent and is defined as follows:

$bs_{Ai} = \{bs_i, Lbs_i, Sbs_i, t_p\}$ where:

$Lbs_{Ai} = \{l_0, ..., l_{359}\}$ denotes the set of petition lists with $l_j = \{p_m, ... p_n\}$ as a single petition list for degree j and p_i is the petition containing information from u_{Ai},

$Sbs_{Ai} = \{s_0, ..., s_n\}$ denotes the list of active sectors at the site with s_k being the k^{th} sector at the site,

$t_p > 0$ denotes the Petition Threshold.

When the Site-Agents receive petitions from User-Agents these petition details, p_{ui}, are stored within an array, in a position determined by the bearing from the site to the user. This allows the site to build up a geographic picture of its petitioning users. When the Site phase starts the first task is to assess the need for a new sector. The Site determines the 90 degree section of its locality that contains the most petitions. If this number exceeds the Petition Threshold, t_p, (which is calculated from the average number of users necessary to make a sector profitable) a sector is created.

Site-Agents are able to configure every aspect of the sectors that they place. The decision making behind the sector configuration is again driven by local knowledge. The next part of the stage involves changing the configuration of sectors. The Site-Agent has information about its petitioners. For example, the site will know that of the petitioners in a certain area where there exists a certain sector, a larger percentage of the petitioners are covered rather than served by this sector. This will lead the Site-Agent to investigate. If there is not enough capacity a new channel will be added. If there are problems with the carrier-to-interference ratio, the site will reject or swap a high interference channel, or change the polarisation of the antenna. If the majority of served users, or petitioners, are in one portion of the sector's antenna coverage area, that antenna may be changed – a simple alteration of azimuth or tilt or a swap to a reduced/ increased beam-width option. Finally, if the change does not meet the sites criteria- a general improvement in the solution or profit level, although backward steps may also occasionally be allowed - it will be undone.

Another aspect of the Site-Agents' behaviour is the co-operation, or competition, between themselves. In a situation where two Site-Agents have sectors attempting to serve the same set of users the network will be fundamentally unstable. This comes from high levels of distraction amongst the User-Agents. As each Site-Agent attempts to satisfy more of the users a 'critical mass' point will be reached and one of the competing sectors will become infeasible. At this point it will be removed, as explained in the next paragraph. However it may occur that two barely viable sectors can co-exist. In this case the User-Agents are aware of the numbers of their neighbours being served from elsewhere and this information is passed to the Site-Agents. Inter-Site-Agent communication takes place, in order to agree to a merging of these sectors. One Site-Agent, usually the one with the most existing infrastructure – and therefore the least relative cost per sector, is handed all the existing users and increases capacity in an attempt to serve them. The other sector is removed from the network.

Site-Agents are able to remove a sector from the network, if that sector is deemed unfeasible or unprofitable, and this is the final consideration of the Site-Agent within each Site-Phase. A sector should be removed if all its users have deserted it, or, over a period of time it is unprofitable and not attracting new users. As the User-Agents swarm from one site to another, over the course of the optimisation, many sectors are created and destroyed. However, Site-Agents are encouraged to be patient with the sectors they create. As with the User-Agent's level of resentment, Site-Agents initially make only smaller, more focused changes.

Preliminary Results

This brief presentation of some results will be replaced in the final submitted paper by a more comprehensive account of the schemes abilities and shortfalls. Using a data set covering the town of Malvern, in the U.K., the agent optimisation system can be seen to be working. Figure 1 shows the progress of the tool over 500 iterations. The graph shows increasing levels of profit (defined as total user revenue minus all infrastructure and running costs), with the final best-profit network, shown in figure 2, making approximately 80000 units profit over a one and a half year return period.

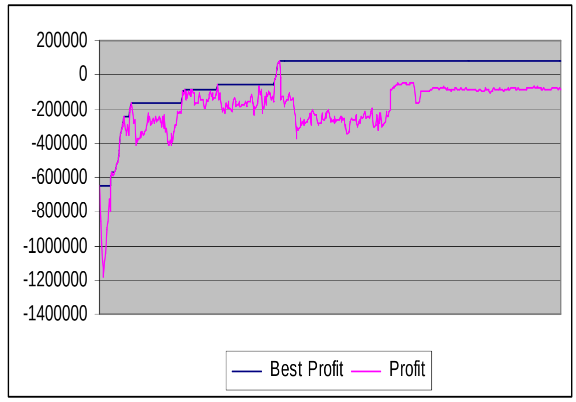

Figure 1. Profit over Iterations

It is interesting to note the schemes preference for the classic four sector pattern (see figure 2). This is not explicitly included in the optimisation but rather a result of the network model. At present, the individual sector optimisation aspect is still preliminary and at present is somewhat ineffective (as illustrated by the very small light green North West sector on the most southerly active site).

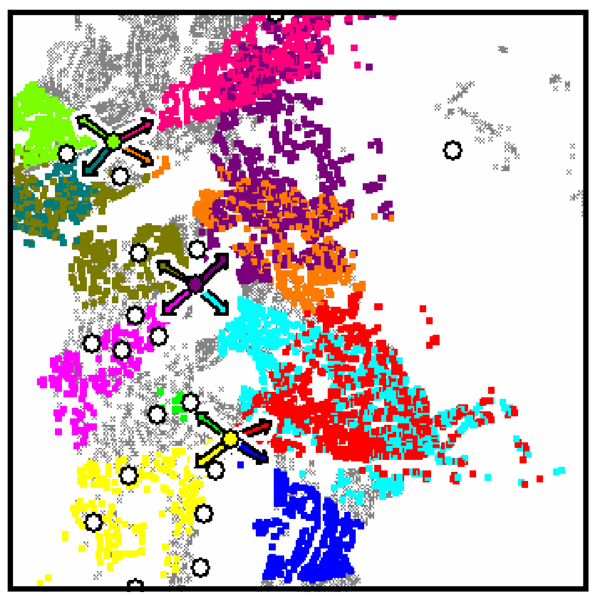

Figure 2. The finished network plan

A poorly optimised sector sharing a site will find it easier to reach profitability because real-world costs such as antenna towers and site maintenance will be shared out among the other sectors at the site. As the sector optimisation code is improved it is expected that some sectors will be alone at a site in best-profit networks.

Figures 3, 4 and 5 illustrate some properties of the best-profit network found by the agent optimisation. The high coverage level is a trademark of the scheme and comes from each individual user's need for service. At this stage in the implementation it can be seen that the optimisation is rather too biased toward the User-Agents high coverage and service demands. Higher levels of service will of course translate into higher revenues. Unfortunately, the Site-Agents, through their actions in the pursuit of higher service levels, are in fact restricting service to large portions of the network.

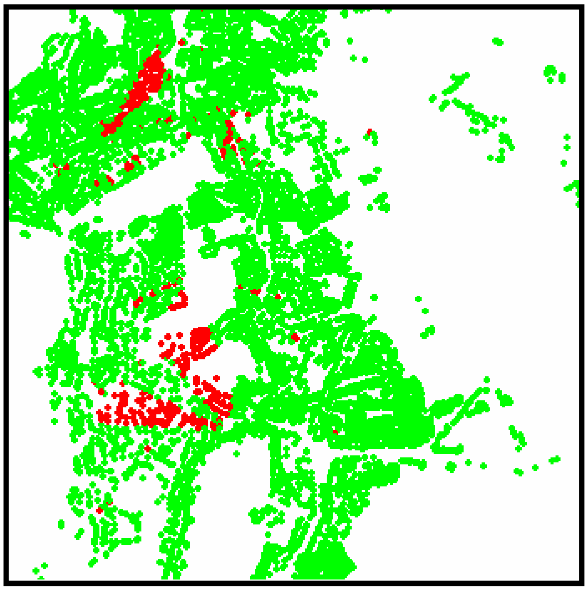

Figure 3. Coverage (in green) at 94%

The fact that this network is in no way limited by capacity is a pointer to this. Indeed almost twice the bandwidth that is actually allocated is supplied by the Site-Agents in an attempt to serve more and more users. This over-engineering manifests as a problem in the high levels of Carrier-to-Interference errors in certain parts of the network (see figure 5), preventing the User-Agents from receiving the very service that the Site-Agents are trying to provide.

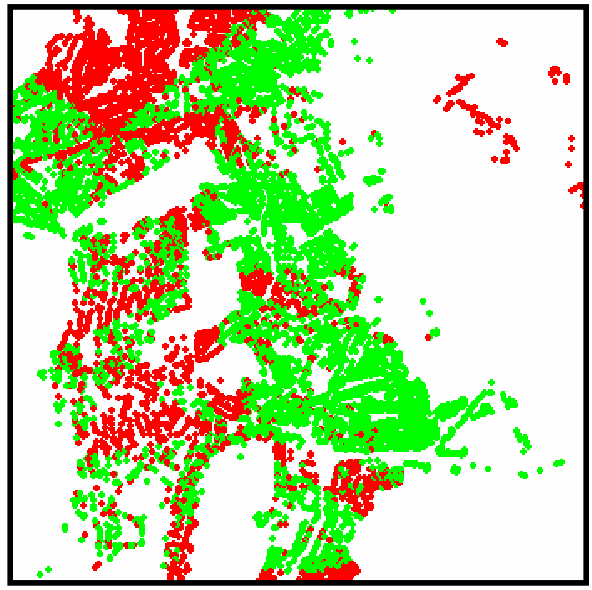

Figure 4. Service (in green) at 63%

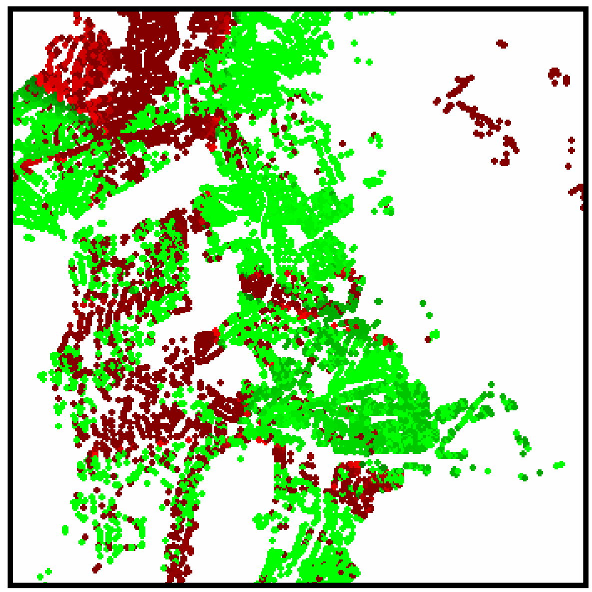

Figure 5. Carrier-To Interference errors (in red)

Conclusions

At the time of writing the implementation of the scheme is progressing well and the results shown here can only hint at the final potential of the approach. However the authors are confident that this scheme will be a viable alternative to the Tabu-Search approach currently implemented in the ECHO tool used for FWA network planning. The approach will firstly replace the existing, time consuming, constraint-satisfaction second stage [5], and eventually should be shown to be comparable to the over-all effectiveness of the Tabu-Search scheme.

References

[1] Tjelta, T., Nordbotten, C., Adams, C., and Craig, K.H. *Millimetre Broadband Radio Access Optimised to Serve the Business as well as the Residential Market: Views from the IST Project Embrace*. in IST Mobile Communications Summit. Galway. 2000.

[2] Glover, F., Taillard, E., and de Warra, D., *A User's Guide to Tabu Search.* Annals of Operations Research, 41: pp. 3-28, 1993.

[3] Allen, S.M., *et al. A Software Tool for Optimised Millimetric Radio Access Cell Planning*. in IST Mobile Communications Summit. Thessaloniki. 2002.

[4] Allen, S., *et al. Maximising Spectrum Efficiency for BFWA through Automatic Cell Planning*. in IEE " Getting the Most Out of the Radio Spectrum" Conference. 2002.

[5] Allen, S.M., Hurley, S., Taplin, R.K., and Wade, A.A. *Cost-Effective Cell Planning of Millimetre Broadband Access Systems*. in ICAP2003. Exeter. 2003.

[6] Dorigo, M., Caro, G.D., and Gambardella, L.M., *Ant Algorithms for Discrete Optimization.* Artifical Life, 5(3): pp. 137-172, 1999.

[7] Socha, K. and Kisiel-Dorohinicki, M., *Agent-based evolutionary multiobjective optimisation*, in Cec'02: Proceedings of the 2002 Congress on Evolutionary Computation, Vols 1 and 2. p. 109-114, 2002.

[8] Dorigo, M., Maniezzo, V., and Colorni, A., *The Ant System: Optimization by a colony of cooperating agents.* IEEE Transactions on Systems, Man, and Cybernetics, 26(1): pp. 1-13, 1996.