Glyndŵr University
Glyndŵr University Research Online

Computing

Computer Science

7-1-2006

Extended End-to-End Cost Metrics for Improved Dynamic Route Calculation

Nigel Houlden
Glyndwr University, n.houlden@glyndwr.ac.uk

Vic Grout
Glyndwr University, v.grout@glyndwr.ac.uk

John McGinn
Glyndwr University, j.mcginn@glyndwr.ac.uk

John N. Davies
Glyndwr University, j.n.davies@glyndwr.ac.uk

Follow this and additional works at: http://epubs.glyndwr.ac.uk/cair

Part of the Computer and Systems Architecture Commons, Digital Communications and Networking Commons, Hardware Systems Commons, and the Systems and Communications Commons

Recommended Citation

Houlden, N., Grout, V., McGinn, J. & Davies, J. (2006) 'Extended End-to-End Cost Metrics for Improved Dynamic Route Calculation', [Paper presented to the 6th International Network Conference (INC 2006) 11th-14th July 2006]. University of Plymouth

This Article is brought to you for free and open access by the Computer Science at Glyndŵr University Research Online. It has been accepted for inclusion in Computing by an authorized administrator of Glyndŵr University Research Online. For more information, please contact d.jepson@glyndwr.ac.uk.

Extended End-to-End Cost Metrics for Improved Dynamic Route Calculation

Abstract

This paper considers the use of compound cost functions in routing calculations. Using an abstracted version of Cisco's EIGRP as its basic model, it develops the theoretical principals of optimal end-to-end interior routing then details the limitations of conventional and current implementation. The requirements of an improved system are discussed and proposals for an enhanced Ant Colony Optimisation - DUAL protocol given. A comparative example is used to illustrate the points made and further work needed and other open questions are considered in conclusion. The paper has two purposes. In the main, it provides an analysis of current routing protocols and a model for future ones. In part, however, it is also intended to promote debate into many aspects of Internet routing and its 'optimality' in advance of long-term development of the new protocol.

Keywords

Routing protocols, Interior routing, Cost metrics, Optimisation, EIGRP

Disciplines

Computer and Systems Architecture | Digital Communications and Networking | Hardware Systems | Systems and Communications

Comments

This paper was presented at 6th International Network Conference (INC 2006) on the 11th-14th July 2006 which was held in Plymouth, UK. Details of the conference are available at http://www.cisnr.org/

This article is available at Glyndŵr University Research Online: http://epubs.glyndwr.ac.uk/cair/58

Extended End-to-End Cost Metrics for Improved Dynamic Route Calculation

Nigel Houlden, Vic Grout, John McGinn and John Davies

Centre for Applied Internet Research (CAIR), University of Wales, NEWI, Wrexham, UK
{n.houlden|v.grout|j.mcginn|j.n.davies}@newi.ac.uk

Abstract

This paper considers the use of compound cost functions in routing calculations. Using an abstracted version of Cisco's EIGRP as its basic model, it develops the theoretical principals of optimal end-to-end interior routing then details the limitations of conventional and current implementation. The requirements of an improved system are discussed and proposals for an enhanced *Ant Colony Optimisation - DUAL* protocol given. A comparative example is used to illustrate the points made and further work needed and other open questions are considered in conclusion. The paper has two purposes. In the main, it provides an analysis of current routing protocols and a model for future ones. In part, however, it is also intended to promote debate into many aspects of Internet routing and its 'optimality' in advance of long-term development of the new protocol.

Keywords

Routing protocols, Interior routing, Cost metrics, Optimisation, EIGRP

1. Introduction: Classes of Routing Protocol and their Performance

Interior routing protocols, through the exchange of topological and other network information, allow path costs to be calculated and routing tables built within routers. There are two classes of protocols with a non-empty intersection.

- *Distance Vector* (*DV*) protocols generally combine the information sharing and route calculation processes. Routers running a purely DV protocol will be aware of the 'distance' to a remote network and the exit interface to best use to reach it. 'Distance' is often a crude cost measure such as hop count (Black, 1999). Such protocols are typically simple but inefficient, slow to converge and distinctly sub-optimal.
- *Link State* (*LS*) protocols exchange status updates to build a topological knowledge of the network. Path calculation is then a separate process and may involve a more sophisticated cost function, based upon different metrics. LS protocols usually give better routing solutions and converge quicker than DV protocols but can have a significant complexity (Doyle, 2005).
- *Hybrid* protocols attempt to combine the desirable features of DV and LS protocols to provide both flexibility and efficiency, minimising routing cost and speeding convergence. Probably the best known example is Cisco's proprietary *Enhanced Interior Gateway Routing Protocol* (*EIGRP*). (Retana *et al.*, 2000).

Almost all forms of practical routing problem are known to have NP-complete complexity (Wang and Crowcroft, 1996). Also, the practical requirements of the network administrator are not always understood by the algorithmic designer (Maltz *et al.*, 2004 and Grout, 2005).

The limitations and sub-optimality of conventional LS routing protocols are discussed by Houlden and Grout (2005). These concerns include the following.

- LS protocols generally use additive end-to-end cost functions, which are inappropriate in the case of certain metrics such as bandwidth or reliability.
- LS protocols optimise paths independently, giving no consideration to how routes *compete* for use of available network connections.

The common notion that LS protocols provide 'optimal' routing is false: hybrid protocols generally provide a more *realistic* (although still imperfect) end-to-end cost calculation. EIGRP's *Diffusing Update Algorithm* (*DUAL*) (Garcia-Luna-Aceves *et al.*, 1999 and Cisco, 2005a), for example, distributes reported path costs in a conventional DV manner but *individually* for different metrics. Moreover, whilst a metric such as delay is appropriately summed, it is the *minimum* bandwidth that is recorded, etc. A significant feature of EIGRP is its ability, through the distribution of both individual metrics and compound cost calculations, to find alternative routes quickly in response to link failures without forming routing loops. However, even these more sophisticated protocols are still less than optimal in other respects. This paper, using (for reasons that will become clear, a *theoretical*) model of EIGRP, outlines the remaining shortcomings and their consequences before suggesting and evaluating possible solutions.

2. A Model for Calculation of End-to-End Path Costs

In this section, the principles of EIGRP are (partially) extended to a general routing model.

2.1. EIGRP

There are various ambiguities and anomalies buried deep in the EIGRP specification (Cisco 2005a & 2005b). An apparently minor example is the impression given that *maximum transmission unit* (*mtu*) can be factored into path costs. In fact, mtu information, along with hop-count, *is* propagated through domain topology tables by EIGRP's *DUAL* but *not* used in the cost formula. mtu is used to set timings between EIGRP packets and hop-counts limit the size of an EIGRP domain. However, a case *could* be made for allowing their consideration in determining the cost of a path.

Another issue arises from considering the path cost function as given. The formula is generally stated as

$$C = \left(k_1 b + \frac{k_2 b}{256 - l} + k_3 d \right) \left(\frac{k_5}{r - k_4} \right) \tag{1}$$

where b is the minimum bandwidth, l the load, d the total delay and r the reliability along the length of the path. k_1, k_2, k_3, k_4 and k_5 are administrator-configurable coefficients (although the values must be consistent across the domain). However, even this calculation is complicated by the need to scale bandwidth and delay as $b = (256 \times 10^8) / b_0$ and $d = 256 d_0$ where b_0 and d_0 are the measured or configured values – the 256 arises from a storage difference (from IGRP to EIGRP) between 24 and 32 bits. b_0 is measured in kbps. Delay is measured in tens of milliseconds on interfaces but used in the formula as milliseconds, requiring an additional scaling by a factor of ten. Integer arithmetic is used in practice, resulting in significant truncation. It is then claimed that the default coefficient values of $k_1=1, k_2=0, k_3=1, k_4=0$ & $k_5=0$ lead to the simplified path cost of $C = b + d$ but this is clearly untrue. If $k_5=0$ then $C = 0$ and all paths will have equal (vanishing) cost. There is a discrepancy between the widespread specification and the implementation. In fact a more detailed reference (Cisco, 2005c) notes that (1) only applies if $k_5>0$. If $k_5=0$, the formula

$$C = \left(k_1 b + \frac{k_2 b}{256 - l} + k_3 d \right) \tag{2}$$

is used instead. Harris and Köhler (1999) make a similar point but misrepresent the relationship between (1) and (2) slightly ($k_4=0$ is *not* necessary for (2) to apply). Also, in fact, if $k_1=k_2=k_3=0$, then a further modification is necessary. As an additional point, it is unclear from any version of the specification how other apparently dynamic metrics such as load and reliability actually propagate. It may well be that EIGRP, in inheriting much of its structure and formularisation from the earlier Cisco *Interior Gateway Routing Protocol* (*IGRP*) (Black, 1999), is now overcomplicated for its purpose and has never been fully implemented within the router Internet Operating System. Finally, the widespread myth that $k_5>0$ factors mtu into the cost calculation is also clearly untrue from considering (1).

2.2. A Generalised End-to-End Routing Model

To simplify and progress the argument, we will abstract the principles of EIGRP into a coherent form, which also includes mtu and hop-count.

Define b_i, d_i, l_i, r_i, and m_i to be the bandwidth, delay, load, reliability (measured as a probability of failure) and mtu respectively of the link i. For a given path P, the end-to-end bandwidth, delay, load, reliability, mtu, and hop-count can be calculated as follows.

$$b_P = \min_{i \in P} b_i \tag{3}$$

$$d_P = \sum_{i \in P} d_i \tag{4}$$

$$l_P = \max_{i \in P} l_i \tag{5}$$

$$r_P = \prod_{i \in P} r_i \tag{6}$$

$$m_P = \min_{i \in P} m_i \tag{7}$$

$$h_P = \left| \{ i : i \in P \} \right| \tag{8}$$

(Minimum bandwidth and mtu and maximum load are propagated. Delays are added and reliabilities - probabilities of failure - multiplied. h is simply the length of P. An alternative is to *measure* metrics such as reliability on an end-to-end basis as EIGRP is often portrayed as doing.) The compound cost of the path P is then some function

$$C(P) = f(b_P, d_P, l_P, r_P, m_P, h_p). \tag{9}$$

In principle, f can then be defined, by reintroducing scaling functions and coefficients as before, as

$$C(P) = k_b f_b(b_P) + k_d f_d(d_P) + k_l f_l(l_P) + k_r f_r(r_P) + k_m f_m(m_P) + k_h f_h(h_P) \tag{10}$$

with f_b, f_d, f_l, f_r, f_m and f_h the required scalings and the coefficients k_b, k_d, k_l, k_r, k_m and k_h set by the administrator. DUAL, or something similar, can distribute and modify the values of b_P, d_P, l_P, r_P, m_P and h_P. The optimum route to any remote network is then found by minimising $C(P)$ for all paths P that lead to it. In EIGRP, this is also a function of DUAL, passing *reported distances* and calculating *feasible distances*.

3. Limitations of the Model and Proposed Extensions

Even this formulation, however, has its shortcomings in implementation. Three are discussed in this section.

3.1. Cost Functions

Firstly, the calculation of $C(P)$ in (10) may be simplistic. It may not be appropriate to sum metrics that are combined other than additively as DUAL propagates routes across the domain. A more sophisticated form for f is probably required but this is far from a simple problem and cannot be considered in a paper of this length. However, before leaving it, it should be noted that the EIGRP formula (1) & (2), though crude, is an attempt to achieve this (Harris and Köhler, 1999). Unfortunately the EIGRP/DUAL *feasibility conditions* (Albrightson *et al.*, 1994) rely on a linear formula! (See next sub-section.)

3.2. Dynamic and Non-Dynamic Metrics

Secondly, there is a common misconception regarding certain metrics. Whilst (minimum) bandwidth, (minimum) mtu and hop-count can reasonably be regarded as fixed for any given path; load, delay and (possibly) reliability should be considered dynamic - they vary with changing traffic flows. In fact, in EIGRP some of these metrics are *fixed* whilst others have dubious foundation. Delay, for example, is a constant derived from the connection type or statically configured by the administrator on the link interface. The false impression (Cisco, 2005a) that there is a dynamic aspect to this process comes from DUAL calculating compound metrics from static values in real time as routes propagate. However, these values are *not* always read dynamically from the network/domain. Ideally, of course, they should be but there is a problem with this notion.

Take load as an example. A high dynamic load over a link or path will discourage routes from using that link/path, which in turn will reduce the traffic, and hence the load - making the link/path favourable once more. The process cycles and routing will be unstable. Some protocols use variants on a *hold-down timer* to damp this effect. A slightly different, but comparable problem is identified by Shaikh *et al.* (2000). Cisco (2005a) discourage the use of k_2, k_4 and k_5 anyway, on the basis that routing loops may be formed although there may in fact be many reasons for this. The application of the purely linear feasibility conditions (Albrightson *et al.*, 1994) for determining feasible successors, for example, do not work for other than the default values and the calculations involving load and reliability are unclear. (Only the values $k_1=1, k_2=0, k_3=1, k_4=0$ & $k_5=0$ give a linear cost function in which path cost can be calculated by adding individual link costs.) This, however, assumes these dynamic factors to be considered purely in *response* to changing network values. The ideal is for them to be *calculated* or *predicted* as discussed in the next sub-section.

3.3. Independent and Combined Route Calculations

Finally, the preferred routes derived by DUAL are, as with their LS counterparts (Houlden and Grout, 2005), calculated independently. In fact, routes *compete* for the use of links and certain metrics will change as links are shared. Load and delay will increase with higher levels of traffic while *available* bandwidth and reliability decrease. Routing an increasing number of traffic streams to use a high bandwidth link, for example, will eventually overload the link. The optimal routing for the domain as a whole will not result from optimising routes individually. *Load balancing*, such as through the use of `variance` in EIGRP (Retana *et al.*, 2000) partially addresses this issue but its implementation is restricted in all established protocols – where it is permitted at all. The instability of dynamic routing factors can, in principle, be overcome by considering the interaction of traffic flows in different routings. This is illustrated by an example in the next section.

4. A Routing Example

For convenience, we use a simple, intuitive model based on load, calculated as traffic flow as a fraction of available bandwidth. $k_b=0, k_d=0, k_l=1, k_r=0, k_m=0, k_h=0$ and $f_l = l_P / b_P$ so that $C(P) = l_P / b_P$ also.

Figure 1 shows the basic topology. Three routers, A, B & C, serve networks X, Y & Z. and are connected by links of different multiples of some arbitrary base bandwidth *ß*. Traffic at a given rate *a* flows from X to Y and from X to Z. (Actual values for *a* and ß in bits per second can be substituted at any point - the intention is show a general principle.) We evaluate two possible routings: AB/ABC, in which traffic from X to Y is carried by AB and traffic from X to Z by ABC (ie, AC via B) and, alternatively, AB/AC, in which XY is carried by AB as before but XZ by AC. Firstly, the routes are considered independently, that is the effects of the combined traffic are ignored. Secondly, these increased levels of traffic are taken into account.

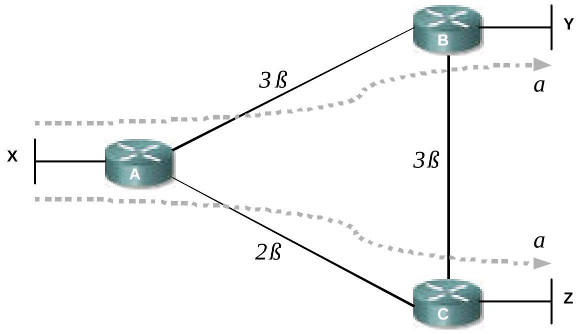

Figure 1. A General Routing Example

4.1. Independent Routes

Consider first the routing AB/ABC. Calculating each route separately, the maximum load on, and therefore cost of, the path XY over AB is $C_{XY(AB)} = a / 3ß$. Similarly, the cost of XZ over ABC is $C_{XZ(ABC)} = a / 3ß$. The cost of the routing AB/ABC is then $C_{AB/ABC} = C_{XY(AB)}+C_{XZ(ABC)} = a / 3ß + a / 3ß = 2a / 3ß$.

Now, for the routing AB/AC, $C_{XY(AB)} = a / 3ß$ as before but $C_{XZ(AC)} = a / 2ß$, giving $C_{AB/AC} = C_{XY(AB)}+C_{XZ(AC)} = a / 3ß + a / 2ß = 5a / 6ß$.

$C_{AB/AC} = 5a / 6ß > 2a / 3ß = C_{AB/ABC}$, making AB/ABC the preferred routing when paths are calculated independently.

4.2. Combined Routes

However, if we consider the combined effects of traffic on the links, we get a different result. For the routing AB/ABC, the combined traffic on the link AB is $2a$ so $C_{XY(AB)} = C_{XZ(ABC)} = 2a / 3ß$ and $C_{AB/ABC} = C_{XY(AB)}+C_{XZ(ABC)} = 2a / 3ß + 2a / 3ß = 4a / 3ß$.

For AB/AC, however, $C_{XY(AB)} = a / 3ß$ and $C_{XZ(AC)} = a / 2ß$ as in the independent case so $C_{AB/AC} = C_{XY(AB)} + C_{XZ(AC)} = a / 3ß + a / 2ß = 5a / 6ß$.

$C_{AB/AC} = 5a / 6ß < 4a / 3ß = C_{AB/ABC}$, making AB/AC the preferred routing when the effects of combined traffic is known. This is the better approach for the good of the domain as a whole. Individually sub-optimal routes combine to give an optimal domain solution.

5. Implementation Issues and Proposals

First, note in the example of the previous section that it is unnecessary to know the exact traffic flows. It will make for better routing strategies if these levels *are* known of course but, in fact, existing routing protocols assume all traffic flows to be equal - that is they give routes equal prominence. With a flexible routing model (Houlden and Grout, 2005), traffic flows may be built in or omitted entirely to fit the application.

Our objectives then are to derive an optimal routing for the entire domain, for a flexible cost function. There are a number of obstacles to this ideal, some partially considered by EIGRP/DUAL, some not.

- There is no current provision for the use of truly dynamic factors such as delay in route cost calculations.
- Internet routing is a distributed process. Although routers share routing information (connectivity and metrics) they do not share routing intent. The form of cooperation necessary to select individually suboptimal routes to achieve routing optimality for the domain implies centralised control.
- Routing is complex (Wang and Crowcroft, 1996) but made manageable by independent route calculations. A truly optimal routing algorithm will be significantly more complex and may overload router processors and/or links between them.

The work of Grout *et al.* (2004) offers some solutions in relation to an OSPF-type link-state routing protocol but cannot deal with dynamic metrics or an EIGRP-type multi-metric cost function. An alternative is proposed here: *ACO-DUAL*.

5.1. The *Ant Colony Optimisation - Diffusing Update Algorithm* (*ACO-DUAL*)

Ant Colony Optimisation (*ACO*) (Dorigo, 2005) is an established approach to large-scale combinatorial problems that has already shown some promise in simulations with fixed costs on networks of restricted size (Johnson and Perez, 2005). An ACO solution consists of small agent packets (*ants*) moving through the problem space – in this case the domain – laying down an electronic *pheromone* and, as a result, propagating information from node to node (router to router). The essential difference between an ant and a routing update is that ants continue through the domain whereas the life of an update is a single hop between routers. ACO-DUAL has three essential components: *Ants*, *Balancing* and *Cost* (*ABC*).

A. Ants

Ant Packets (*AP*s) will perform the equivalent EIGRP functions of both ‘Hello’ and ‘Update’ packets. An ACO-DUAL-enabled router will send an AP, consisting of metrics b_P, d_P, l_P, r_P, m_P and h_P, to each ACO-DUAL neighbour, which will recalculate metrics as given in equations (3)–(8) and calculate its compound cost as in (C.). The AP is then updated and copied to all neighbours and the process continues as APs propagate across the domain. Where APs from the same remote network, via different routes, have converged at a router, all but the AP of lowest calculated cost are redundant and only this AP is forwarded (see

Figure 2). Load, delay and reliability are taken (truly) *dynamically* from the interface at each point. The EIGRP use of reported distance to install backup routes is avoided by dynamic cost calculation (C.) and load balancing (B.)

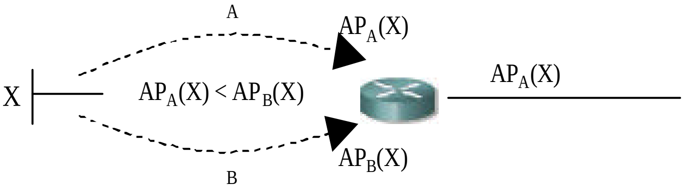

Figure 2. Ant forwarding

B. Balancing

ACO-DUAL will enforce *auto-variance*. An ACO-DUAL router, having calculated its cost to a remote network via each route (interface) will *automatically* install all routes up to a given factor above the cheapest into the routing table. Packets to each remote network will be balanced across each available route (interface) in proportion to these costs. This automatic load-balancing in response to truly dynamic metrics should be stable and thus avoid the use of hold-down timers.

C. Cost

The ACO-DUAL cost function, $C(P) = f(b_P, d_P, l_P, r_P, m_P, h_p)$, will be *fixed*. Not only will this remove the EIGRP ambiguities of sub-section 2.1., it will, in the correct form, make the routing inherently stable yet responsive to changes without forming routing loops. Costs are retained within routers: reported costs are *not* propagated. There is no need for this as the use of auto-variance (B.) automatically provides for backup routes.

6. Conclusions and Future Work

The *intentions* of DUAL are sound. Propagating individual performance metrics across the domain rather than crude or combined distance measures gives greater routing ‘awareness’. This combined with EIGRP’s load-distribution should produce a routing strategy considerably nearer to domain-optimality than would be possible with purely DV or LS protocols. The problem is with the *implementation* and may even be in part historical. In practice, load-distribution is limited, supposedly dynamic metrics are actually static, the generalised EIGRP cost function does not support the DUAL operation and there is a great deal of uncertainty regarding some other components and doubt as to whether all have even been implemented!

Having identified these shortcomings, this paper has put in place a model for a truly dynamic, diffusing routing protocol and demonstrated its value through a simple example. Comments are sought on these specific proposals for ACO-DUAL. A large-scale simulation programme is about to begin on the ns-2 network simulator (ns2, 2005) in order to establish the proposed protocol in principle and consider the following two issues in particular. Firstly, the intention is to determine the best (most accurate and stable) form of $C(P) = f(b_P, d_P, l_P, r_P, m_P, h_p)$. Secondly, the ideal of infinite variance for load-balancing purposes must be examined for practicality. If there are combinatorial (complexity) obstacles to this then it may be necessary to limit the number of alternate routes in the routing table. It is anticipated that balancing will still be automatic and in proportion. An alternative may be to transmit APs only on the higher bandwidth interfaces.

7. References

Albrightson, B., Garcia-Luna-Aceves, J.J. and Boyle, J. (1994) "EIGRP - A Fast Routing Protocol Based on Distance Vectors", *Proceedings of. Networld/Interop 94*, Las Vegas, Nevada, May 1994.

Cisco (2005a) "Enhanced Interior Gateway Routing Protocol", Document ID: 16406, http://www.cisco.com/en/US/tech/tk365/technologies_white_paper09186a0080094cb7.shtml (Accessed 5 September 2005)

Cisco (2005b) "Setting a Preferred Route by Influencing EIGRP Metrics", Document ID: 13673, http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a00800c2d96.shtml (Accessed 5 November 2005)

Cisco (2005c) "IP Routing Protocols Commands: K through M", http://www.cisco.com/en/US/products/sw/iosswrel/ps5187/products_command_reference_chapter09186a008017d029.html#wp1040267 (Accessed 5 September 2005)

Dorigo, M. (2005) "Ant Colony Optimization", http://iridia.ulb.ac.be/~mdorigo/ACO/ACO.html (Accessed 24 November 2005)

Doyle, J. (2005) *OSPF and IS-IS: Choosing an IGP for Large-Scale Networks*, Addison Wesley.

Garcia-Luna-Aceves, J.J., Vutukury, S. and Zaumen, W.T. (1999) "A Practical Approach to Minimizing Delays in Internet Routing", *Proceedings of the IEEE International Conference on Communications (ICC'99)*, Vol. 1, pp479-483.

Grout, V. (2005) "Ne'er the Twain Shall Meet? Bridging the network optimisation reality gap", *Mathematics Today*, Vol. 41, No. 5, October 2005, pp157-160.

Grout, V., Davies, J., Hughes, M. & Houlden, N. (2004) "A New Distributed Link-State Routing Protocol with Enhanced Traffic Load Distribution", *Proceedings of the 4th International Network Conference (INC 2004)*, University of Plymouth, 6th-9th July 2004, pp165-172.

Harris, R. and Köhler, S. (1999) "*Possibilities for QoS in Existing Internet Routing Protocols*", Report No. 228, University of Würzburg Institute of Computer Science, Germany.

Houlden, N. and Grout, V. (2005) "Principles of Optimal Interior Routing", *Proceedings of the 1st International Conference on Internet Technologies and Applications (ITA 05)*, Wrexham, North Wales, UK, 7th-9th September 2005, pp274-281.

Johnson, C.M. and Perez, E. (2005) "An Ant Colony Optimization Algorithm for Dynamic, Multi-Objective Network Routing", *Proceedings of the International Conference on Internet Technologies and Applications (ITA 05)*, 7th-9th September 2005, Wrexham, UK, pp293-300.

Maltz, D.A., Xie, G., Zhan, J., Zhang, H., Hjálmtýsson, G. and Greenberg, A. (2004) "Routing Design in Operational Networks: A look from the inside" *Proceedings of ACM SIGCOMM'04*, 30th August-3rd September 2004, Portland, Oregon, USA, pp27-40.

ns2, (2002), "*The Network Simulator – ns – 2*", www.isi.edu/nsnam/ns (Accessed 15 November 2005)

Retana, A., White, R. and Slice, D. (2000) *EIGRP for IP: Basic Operation and Configuration*, Addison Wesley.

Shaikh, A., Kalampoukas, L., Dube, R. and Varma, A. (2000) "Routing Stability in Conjested Networks: Experimentation and analysis", *Proceedings of ACM SIGCOMM'00*, Stockholm, Sweden, 2004, pp163-174.

Wang, Z. and Crowcroft, J. (1996) "Quality of Service Routing for Supporting Multimedia Applications", *IEEE Journal on Selected Areas in Communications*, Vol. 14, No. 7, September 1996, pp1228-1234.