Vic Grout, Nigel Houlden, John Davies, John McGinn and Stuart Cunningham

Centre for Applied Internet Research, University of Wales, NEWI Plas Coch
Campus, Mold Road, Wrexham, LL11 2AW, UK

A UNIFIED FRAMEWORK FOR OPTIMAL ROUTING

1. Introduction

The complexity and limitations of performance and interoperability of different Internet routing protocols are serious obstacles to future network integration. Global and European Internet traffic is increasing exponentially and will continue to do so with the introduction of new services now and over the next few years. The significance of routing protocols and the importance of their level of optimality grows in proportion. Also the current trend toward political and industrial integration requires different networks to interoperate seamlessly in a manner not considered when their original standards and protocols were defined. An efficient method of discussing, comparing and improving routing and routing protocols, allowing for appropriate route redistribution, for example, is long overdue.

It requires only limited analysis to recognize that existing Internet routing protocols are sub-optimal in their performance. However, any attempt to address the problem is hampered by differences in principles and objectives of the different systems and their underlying algorithms. This paper provides a framework for appropriately comparing and contrasting routing protocols. It has three main parts. Firstly, existing and proposed routing protocols are classified according to their method of operation, scope and objective. This is a significant improvement on the current distinction between distance-vector and link-state protocols. Secondly, a universal theory of path-, network- and domain-optimal routing is developed. This takes into account differences in underlying metrics, cost calculations and objective functions, and is sufficiently general to include all existing protocols and a number of proposed extensions. Finally, with these results in place, the paper describes the shortcomings of existing routing protocols and discusses options for seeking improvement. This covers key aspects of computational complexity and outlines the difference between exact and heuristic optimization. New advances using local search algorithms are introduced and the value of *Ant Colony Optimization* (*ACO*) techniques is assessed in conclusion.

2. A Classification of Routing Protocols

An *internet* comprises a number of *networks* organised into connected *domains*. A domain, sometimes called an *Autonomous System* (*AS*), is a group of networks, usually under a common administration, together with the devices connecting them. These connecting *routers* run *routing protocols* (*RPs*) to calculate preferred routes to remote networks. *Interior Routing Protocols* (*IRPs*) exchange routing information among routers within a domain, *Exterior Routing Protocols* (*ERPs*) between domains. Figure 1 illustrates the structure. ERPs are generally more concerned with applying rules than calculating routes [1] and are not considered here. This paper deals with IRPs.

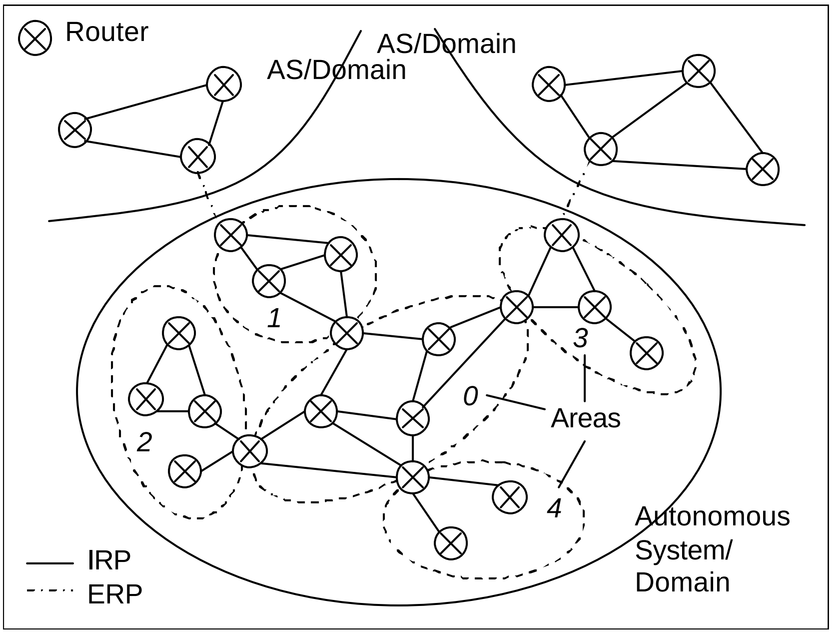

Figure 1: The routing model

IRPs are often classified as *distance-vector* (*DV*) or *link-state* (*LS*) according to the method of information exchange used. A more appropriate classification for the purposes of this paper will be on the basis of the route-calculation algorithms applied. Although there are some aspects of DV protocols that restrict the sophistication of the route calculation [1], much is convention only and the two are treated independently here. Also, where a distinction is necessary, the term *path* will be used to represent the sequence of links from source to destination network and the term *route* for the first such step, this being the entry in the router's *routing table*. A *routing* is the derivation of one or more routes, which may or may not involve knowledge of a complete path.

Table 1. Types of Routing Protocol

	P0	**P1**	**P2**	**P3**	**P4**
Examples:	*Static*	*RIP*	*OSPF*	*EIGRP*	*ERP*
The RP can calculate path costs	✗	✓	✓	✓	✓
The RP can apply different costs to different links	✗	✗	✓	✓	✓
The RP can apply dynamic links costs	✗	✗	✗	✓	✓
The RP can apply cumulative link costs	✗	✗	✗	✗	✓

Five classes of protocol, of increasing sophistication, are defined below and summarised in Table 1.

Type 0: Costless Routing (P0)

A P0 protocol has no concept of link cost, essentially because any route calculation process is suspended or restricted for one or more routes. Transmission over such routes is a fixed, random or broadcast process. Static routes assigned manually by a network administrator trivially use a P0 protocol.

Type 1: Unit-Cost Routing (P1)

A P1 protocol assigns an equal (unit) cost to all links. A P1 protocol will calculate the preferred route as the path with the fewest links. *Routing Information Protocol* (*RIP*) [1] is an example of a P1 protocol.

Type 2: Metric-Cost Routing (P2)

A P2 protocol can accommodate links of different costs. An optimal route will be the path for which the sum of the individual link costs is minimal. For example, the *Open Shortest Path First* (*OSPF*) protocol [2], in its simplest form, calculates cost as 10^8 / b where b is the bandwidth in bits per second of the link [1][3].

Type 3: Dynamic-Cost Routing (P3)

A P3 protocol allows link costs to vary dynamically, reflecting changing network characteristics. As an example Cisco's *Enhanced Interior Gateway Routing Protocol* (*EIGRP*) [1] can be configured to use transient factors such a load or delay to calculate cost. For any given cost configuration (*snapshot*), routes are calculated by minimising path costs as with P2 protocols.

Type 4: Cumulative-Cost Routing (P4)

A P4 protocol recognises that true link costs, dynamic or otherwise, cannot be determined in isolation. A link which forms part of the path for a number of routes may have a higher cost than will be apparent for each route independently (since it will be subject to a greater load or experiencing a longer delay, etc.). Its true cost will be a function of the routing across the whole domain, not for a network or network pair considered separately. P4 protocols, and their performance relative to non-P4 protocols, are discussed in this paper. Only P4 protocols can achieve or closely approximate global routing optimality. Proposals for such an *Enhanced Routing Protocol* (*ERP*) are given in [4] and [5].

3. A Framework for Optimal Routing

In what follows, the notation $i \to j$ is used to represent the single link from i to j and $a \Rightarrow b$ for the path between end points a and b. $a \Rightarrow b \triangleright i \to j$ means that traffic from a to b is carried by the link $i \to j$. $\forall$ is used as shorthand for *'for all'* or *'for every'* and $\exists$ for *'there is'* or *'there exists'*.

Define a domain $D = (N, T)$ by a set of n networks $N = \{1,2,..,n\}$ and a *traffic matrix* $T = (t_{ab}: a,b \in N)$ where t_{ab} represents the traffic (known flow, projected requirement, etc.) from a to b. (In situations in which traffic cannot be measured or predicted, we can set $T = (\underline{1})$, that is $t_{ab} = 1 \ \forall \ a,b \in N$.)

A *protocol P = (M, c),* acting on a domain *D,* is defined by a *metric matrix M* = *(m_{ij}: i,j Î N)* and a *cost function c(t,m)*. m_{ij} specifies the measure of *i*®*j* used by *P* and *c(t,m)* the cost of carrying traffic *t* on a link of metric *m.*

A *distribution X = (x_{ij}^{ab} : a,b,i,j Î N),* acting on a domain *D,* is defined as

$$x_{ij}^{ab} = \begin{cases} 1 : a \Rightarrow b \triangleright i \rightarrow j \\ 0 : otherwise \end{cases} . \qquad (1)$$

A distribution is *feasible* if a sequence of links carries traffic from source destination where necessary. A distribution is *path-feasible* (or *a*Þ*b-feasible*) for *a*Þ*b* if x_{ab}^{ab} = *1* or *$ k_1,k_2,..,k_d* such that $x_{ak_1}^{ab} = x_{k_1k_2}^{ab} = .. = x_{k_{d-1}k_d}^{ab} = x_{k_db}^{ab} = 1$. A distribution is *network-feasible* (or *a-feasible*) for *a* if *" b Î N,* x_{ab}^{ab} *= 1* or *$ k_1,k_2,..,k_d* such that $x_{ak_1}^{ab} = x_{k_1k_2}^{ab} = .. = x_{k_{d-1}k_d}^{ab} = x_{k_db}^{ab} = 1$. A distribution is *domain-feasible* (or *fully-feasible*) if *" a,b Î N,* x_{ab}^{ab} = *1* or *$ k_1,k_2,..,k_d* such that $x_{ak_1}^{ab} = x_{k_1k_2}^{ab} = .. = x_{k_{d-1}k_d}^{ab} = x_{k_db}^{ab} = 1$. In each case, *d* is the *degree of rerouting* for *a*Þ*b* under *X*. (If x_{ab}^{ab} *= 1, d = 0*.)

For a given distribution, *X,* the total traffic on *i*®*j* is given by

$$t_{ij}^{X} = \sum_{a\in N}\sum_{b\in N} x_{ij}^{ab} t_{ab} \qquad (2)$$

and the cost of *i*®*j* by

$$k_{ij}^{X} = c(t_{ij}^{X}, m_{ij}) = c(\sum_{a\in N}\sum_{b\in N} x_{ij}^{ab} t_{ab}, m_{ij}) . \qquad (3)$$

The ideal *routing* is then a fully-feasible distribution *X* that minimises

$$K^{X} = \sum_{i\in N}\sum_{j\in N} k_{ij}^{X} = \sum_{i\in N}\sum_{j\in N} c(t_{ij}^{X}, m_{ij}) = \sum_{i\in N}\sum_{j\in N} c(\sum_{a\in N}\sum_{b\in N} x_{ij}^{ab} t_{ab}, m_{ij}) \qquad (4)$$

across the domain. However, such optimality may be difficult to achieve. It is considered now at three levels.

Define a *path-routing* P^{ab} = *(*p_{ij}^{ab} *: i,j Î N)* for *a Þ b* as $p_{ij}^{ab} = x_{ij}^{ab}$ *" i,j Î N.* P^{ab} is feasible if *X* is *a Þ b*-feasible. Define a *network-routing* Q^{a} = *(*q_{ij}^{ab} *: b,i,j Î N)* for *a* as $q_{ij}^{ab} = x_{ij}^{ab}$ *" b,i,j Î N.* Q^{a} is feasible if *X* is *a*-feasible. Define a *domain-routing* *R* = *(*r_{ij}^{ab} *: a,b,i,j Î N)* as $r_{ij}^{ab} = x_{ij}^{ab}$ *" a,b,i,j Î N.* *R* is feasible if *X* is domain-feasible.

The (known) cost of *i ® j* (solely) under a path-routing P^{ab} is $p_{ij}^{ab} c(t_{ab}, m_{ij})$. The *path-cost* of P^{ab} is then given by

$$C^{ab} = \sum_{i \in N} \sum_{j \in N} p_{ij}^{ab} c(t_{ab}, m_{ij}) = \sum_{i \in N} \sum_{j \in N} c(x_{ij}^{ab} t_{ab}, m_{ij}) \tag{5}$$

(assuming *c(0,m) = 0*). If P^{ab} minimises C^{ab}, P^{ab} is said to be *path-optimal* for *a Þ b*. *X* is path-optimal if P^{ab} minimises C^{ab} *" a,b Î N.* If *X* is path-optimal then

$$K^{X}_{<path>} = \sum_{a \in N} \sum_{b \in N} \sum_{i \in N} \sum_{j \in N} c(x_{ij}^{ab} t_{ab}, m_{ij}) \tag{6}$$

is minimised.

The (known) traffic on *i ® j* under a network-routing Q^{a} is $\sum_{b \in N} q_{ij}^{ab} t_{ab}$ and its cost given by $c(\sum_{b \in N} q_{ij}^{ab} t_{ab}, m_{ij})$. The *network-cost* of Q^{a} is then

$$C^{a} = \sum_{i \in N} \sum_{j \in N} c(\sum_{b \in N} q_{ij}^{ab} t_{ab}, m_{ij}) = \sum_{i \in N} \sum_{j \in N} c(\sum_{b \in N} x_{ij}^{ab} t_{ab}, m_{ij}) . \tag{7}$$

If Q^{a} minimises C^{a}, Q^{a} is said to be *network-optimal* for *a.* *X* is network-optimal if Q^{a} minimises C^{a} *" b Î N.* If *X* is network-optimal then

$$K^{X}_{<network>} = \sum_{a \in N} \sum_{i \in N} \sum_{j \in N} c(\sum_{b \in N} x_{ij}^{ab} t_{ab}, m_{ij}) \tag{8}$$

is minimised.

The traffic on $i \rightarrow j$ under a domain-routing R is $\sum_{a \in N} \sum_{b \in N} r_{ij}^{ab} t_{ab}$ and its cost given by $c(\sum_{a \in N} \sum_{b \in N} r_{ij}^{ab} t_{ab}, m_{ij})$. The *domain-cost* of R is then

$$C = \sum_{i \in N} \sum_{j \in N} c(\sum_{a \in N} \sum_{b \in N} r_{ij}^{ab} t_{ab}, m_{ij}) = \sum_{i \in N} \sum_{j \in N} c(\sum_{a \in N} \sum_{b \in N} x_{ij}^{ab} t_{ab}, m_{ij}). \qquad (9)$$

If R minimises C, R is said to be *domain-optimal.* $X\,(=R)$ is domain-optimal if R is domain-optimal. If X is domain-optimal then

$$K^{X}_{<domain>} (= C) = \sum_{i \in N} \sum_{j \in N} c(\sum_{a \in N} \sum_{b \in N} x_{ij}^{ab} t_{ab}, m_{ij}) \qquad (10)$$

is minimised, which is the derivation of K^X (now $K^{X}_{<domain>}$) in (2), (3) and (4). The ideal then is domain-optimality - that is, to minimise $K^{X}_{<domain>}$. Most existing protocols are path-optimal: they minimise $K^{X}_{<path>}$ for each path $a \Rightarrow b$ independently, which will not, in general, minimise $K^{X}_{<network>}$ for each network a or $K^{X}_{<domain>}$ for the complete domain. Pursuing network-optimality, that is minimising $K^{X}_{<network>}$ independently for each a, will not minimise $K^{X}_{<domain>}$ but will achieve a lower $K^{X}_{<domain>}$ than minimising $K^{X}_{<path>}$ independently for each $a \Rightarrow b$. Using the classification of the previous section, P0 protocols are not capable of any cost minimisation. P1, P2 & P3 protocols will minimise $K^{X}_{<path>}$ with increasingly sophisticated metrics $M = (m_{ij}\!: i,j \in N)$ and costs $c(t,m)$. Only P4 protocols can minimise or reduce $K^{X}_{<network>}$ or $K^{X}_{<domain>}$. Before continuing, we refine the definition of P4 protocols.

Type 4a: Network-Cumulative-Cost Routing (P4a)

A P4a protocol seeks to minimize or reduce

$$K^{X}_{<network>} = \sum_{a \in N} \sum_{i \in N} \sum_{j \in N} c(\sum_{b \in N} x_{ij}^{ab} t_{ab}, m_{ij}) \qquad (11)$$

for each network a.

Type 4b: Domain-Cumulative-Cost Routing (P4b)

A P4b protocol seeks to minimize or reduce

$$K^{X}_{<domain>} = \sum_{i \in N} \sum_{j \in N} c(\sum_{a \in N} \sum_{b \in N} x^{ab}_{ij} t_{ab}, m_{ij}) . \qquad (12)$$

4. Achieving and Approximating Optimality

P0, P1, P2 and P3 protocols have been discussed in length in many forms (eg [1][2][3]). This paper concludes by considering options and possibilities for P4 (a and b) protocols. There are two significant difficulties in the implementation of both types:

- Practical routing protocols cannot employ centralized algorithms – that is, a single process running on some central site. They must run in distributed form – that is, independently on each router. The problem with such independent (P1, P2, P3) processes to date is that they have competed for, rather than co-operated over, network resources [4].
- The complexity of the algorithms becomes an issue. minimizing or reducing $K^{X}_{<domain>}$ is computationally more expensive than minimizing or reducing $K^{X}_{<network>}$, which is in turn a harder process than minimizing $K^{X}_{<path>}$ [5].

Considering the second point, if exact solutions are not achievable then approximate (*heuristic*) methods will be required. Whilst P1, P2 and P3 protocols can minimize $K^{X}_{<path>}$ exactly [6], heuristic methods will almost certainly be required for both P4a and P4b protocols [5].

Two heuristic, distributed (router-based) P4a algorithms are given in [4] and [5]. Both are intended to provide enhanced performance: the first (P4aI in Tables 2 and 3) in comparison with manual domain partitioning (not discussed in this paper) and the second (P4aII) over simple pairwise (P2/P3) network routing. They use novel implementations of established algorithms [7][8] together with the addition of co-operative features. This is essentially possible because $K^{X}_{<network>}$ can be minimised or reduced independently for each network a. Testing has been carried out on two platforms:

- The ns2 network simulator [9]
- The University of Wales NEWI (UoWNEWI) NetSim package

To illustrate the value of such an approach, domains of 20, 50 and 100 networks were tested, 12 runs in each case, with topologies and bandwidths generated at random. Timings are given for the partitioning algorithm, P4aI and the path determination algorithm, P4aII, and in comparison with the equivalent P2 or P3 protocol. Routing cost results are given for a compound algorithm, running independently on each router (as it would in practice) in comparison with a P2/P3 protocol.

Table 2: Run times

Number of networks (m)	(Mean) number of steps to converge			% increase over P3	
	P3	P4aI	P4aII	P4aI	P4aI+P4aII
20	7.2×10^6	8.1×10^5	8.4×10^6	17	28
50	7.4×10^7	8.2×10^6	1.0×10^8	35	47
100	4.3×10^8	4.4×10^7	7.8×10^8	81	92

Table 2 summarizes run times, measured in *steps*. A step is a single high-level language instruction. Convergence is the state of stable partitioning and/or routing across the domain.

The percentage increase in run time of P4aI and P4aII algorithms over P2/P3 protocols increases approximately linearly with the size of the domain.

Table 3 compares $K^X_{<network>}$ summed over all networks a, and $K^X_{<domain>}$ for P2/P3 protocols with P4aI and P4aII algorithms. Although P4aI and P4aII algorithms seek only to improve $K^X_{<network>}$ independently for each network, savings in $K^X_{<domain>}$ also result.

Both algorithms (in combination) offer significant improvements in routing/loading efficiency at the expense of a modest (in complexity terms) increase in run time. P4a protocols then are clearly viable – although they remain to be implemented on production routers in fully operative environments. Whether these heuristic algorithms can be replaced by exact ones is an open question.

Table 3: Routing costs

Number of networks (m)	(total) $K^{X}_{<network>}$			$K^{X}_{<domain>}$		
	% improvement of P4aI+P4aII over P3					
	Minimum	**Maximum**	**Mean**	**Minimum**	**Maximum**	**Mean**
20	14	69	37	5	65	20
50	23	88	43	19	69	32
100	26	72	44	25	75	37

So what of P4b protocols? How can independent processes combine to minimise or reduce $K^{X}_{<domain>}$? The solution may come from an agent-based approach or, more specifically, through the use of ant-colony optimisation (ACO) [10].

ACO algorithms mimic the behaviour of co-operative ants in laying down (sharing) key information (pheremone) to determine a best or improved collective approach to route-finding. These 'ants', in a routing environment, could be used to share routing information, or more specifically routing *intention*, between routers and networks to improve routing across the whole domain. Although heuristic by definition, it may be that significantly better results can be obtained than in isolation.

Success, on a limited scale, in using ant-colony algorithms for adaptive network routing solutions is reported in [11]. Although untested for larger domains, it may be expected that these techniques, in conjunction with the partitioning and path determination concepts in [4] and [5], could lead to further advances. Research continues into ACO network/domain routing in the *Centre for Applied Internet Research* (*CAIR*) [12] at the University of Wales, NEWI Wrexham in the UK.

11

Literature

[1] Slattery, T. and Burton, B. (2000), Advanced IP Routing in Cisco Networks, Osborne McGraw-Hill.

[2] Moy, J.T., (1998), “OSPF: Anatomy of an Internet Routing Protocol”, Addison Wesley.

[3] Moy, J.T., (2000), OSPF Complete Implementation, Addison Wesley.

[4] Grout, V., (2004), “A Self-Partitioning Link-State Routing Protocol”, Proceedings of IEE/IEEE ICN’04, Gosier, Guadeloupe, French Caribbean, 1st-4th March (to appear).

[5] Grout, V., Davies, J., Hughes, M. and Houlden, N. (2004) A New Distributed Link-State Routing Protocol with Enhanced Traffic Load Distribution, Proceedings of BCS/IEE International Network Conference - INC 2004, University of Plymouth, 6th-9th July 2004, pp165-172.

[6] Dijkstra, E.W., (1959), “A Note on Two Problems in Connexion with Graphs”, Numerische Mathematik, Vol. 1, pp269-271.

[7] Prim, R.C., (1957), “Shortest Connection Networks and Some Generalizations”, Bell System Technical Journal, Vol. 36, pp1389-1401.

[8] Kruskal, J.B., (1956), “On the Shortest Spanning Subtree of a Graph and the Traveling Salesman Problem”, Proceedings of the American Mathematical Society, Vol. 7, pp48-50.

[9] ns2, (2002), “The Network Simulator – ns – 2”, http://www.isi.edu/nsnam/ns

[10] Dorigo, M. and Stutzle, T. (2004) Ant Colony Optimization, Bradford Books/M.I.T. Press.

[11] Legge, D. and Baxendale, P., (2003), “An Agent-Managed Ant-Based Network Control System”, Centre for Telecommunication Networks, School of Engineering, University of Durham, UK, http://www.dur.ac.uk/telecoms.networks/public/pdfs/AAMAS-III.pdf

[12] CAIR, (2004) Centre for Applied Internet Research, University of Wales, NEWI Wrexham, http://www.newi.ac.uk/computing/cair

SUMMARY

This paper has defined and analysed different classes (P0, P1, P2, P3, P4a and P4b) of routing protocol. P0, P1, P2 and P3 already exist and are used on production routers. P4a protocols are shown to be viable and worthwhile and possible avenues for P4b protocols are discussed in conclusion.