

A Potts Neuron Approach to Communication Routing

Jari Häkkinen
Martin Lagerholm
Carsten Peterson
Bo Söderberg

*Complex Systems Group, Department of Theoretical Physics, University of Lund,
SE-223 62 Lund, Sweden*

A feedback neural network approach to communication routing problems is developed, with emphasis on multiple shortest path problems, with several requests for transmissions between distinct start and end nodes. The basic ingredients are a set of Potts neurons for each request, with interactions designed to minimize path lengths and prevent overloading of network arcs. The topological nature of the problem is conveniently handled using a propagator matrix approach. Although the constraints are global, the algorithmic steps are based entirely on local information, facilitating distributed implementations. In the polynomially solvable single-request case, the approach reduces to a fuzzy version of the Bellman-Ford algorithm. The method is evaluated for synthetic problems of varying sizes and load levels, by comparing to exact solutions from a branch-and-bound method, or to approximate solutions from a simple heuristic. With very few exceptions, the Potts approach gives high-quality legal solutions. The computational demand scales merely as the product of the numbers of requests, nodes, and arcs.

1 Introduction ---

Communication-routing resource allocation problems are becoming increasingly relevant given the upsurge in demand of the Internet and other telecommunication services. One such problem amounts to assigning arcs in a connected network to requests from start to end nodes, given capacity constraints on the arcs such that a total additive cost (path length) is minimized. (For a review of notation and existing routing techniques, see Bertsekas & Gallager, 1987.) A relatively simple routing problem, with only one request at a time, is the shortest path problem (SPP), which can be solved exactly in polynomial time using, for example, the Bellman-Ford (BF) algorithm (Bellman, 1958) (see also Bertsekas & Gallager, 1987). The multiple shortest path problem (MSPP), where arcs are allocated simultaneously to several re-

quests, is more difficult. We strongly suspect MSSP to be NP-hard, although to our knowledge this has not been proved in the literature.

In this article we address the MSPP using feedback Potts neural networks, which have proved to be powerful in other resource allocation problems, with (Lagerholm, Peterson, & Söderberg, 1997) or without (Gislén, Peterson, & Söderberg, 1992) a nontrivial topology. For each request, we assign a Potts network, with units encoding which arcs are to be used by that request. Appropriate energy terms are constructed in terms of the Potts neurons to minimize total path lengths and ensure that capacity constraints are not violated. Mean field (MF) equations are iterated using annealing to minimize the total energy. In contrast to earlier usage of Potts encoding and MF annealing (Peterson & Söderberg, 1989; Gislén et al., 1992; Lagerholm et al., 1997), where global objective functions are minimized, here each node minimizes its own local energy.

For the case of a single request, the Potts MF approach reduces in the zero temperature limit to the BF algorithm; hence our approach contains this standard algorithm as a special case.

For each request, the Potts MF network (Peterson & Söderberg, 1989) defines an (inverted) “fuzzy” spanning tree (a subgraph without loops connecting all nodes) rooted at the end node. In order to project out the part defining the (fuzzy) path from the start node, and to keep track of the paths in general, we utilize a propagator matrix formalism (following Lagerholm et al., 1997). The computation of the propagator requires matrix inversion; fortunately this can be done using an iterative procedure with a low computational cost. As in the previously considered airline crew scheduling problem (Lagerholm et al., 1997), proper preprocessing is employed to identify independent subproblems in order to reduce the problem complexity.

Despite the existence of global constraints, the implementation of the approach is truly local. When updating the MF equations for a particular node, only information residing at neighboring nodes is needed.

The approach is gauged by an exact branch-and-bound (BB) algorithm (for smaller problems) and by two BF-inspired heuristics (for larger problems), on a set of synthetic but challenging test problems, showing an excellent performance of the Potts MF approach, with a CPU consumption per request scaling merely like NN_L , where N is the number of nodes and N_L the number of links in the network. The method is also very robust with respect to parameters.

Another novel method (Boyan & Littman, 1994) has been proposed for the MSSP. In contrast to our approach, it is aimed at dynamical problems. It is also rooted in the BF algorithm, but is in the static limit unrelated to our Potts algorithm. However, in this limit it reduces to the independent BF approach, which is used for comparisons in this work.

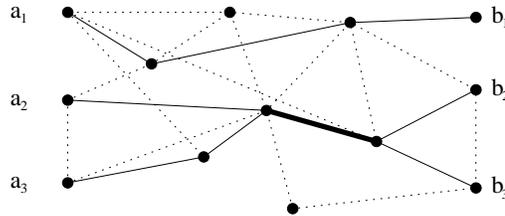


Figure 1: Example of a solution to a three-request problem. Dotted lines represent unused links, and solid lines represent links used by the requests.

2 The Multiple Shortest Path Problem

An MSPP is defined by specifying the following:

- A connected network of N nodes and N_L (bidirectional) links, corresponding to $2N_L$ arcs.
- For each arc ij , a cost (arc length) d_{ij} and a capacity C_{ij} .
- A set of N_R transmission requests r , each defined by a start node a_r and an end node b_r .

The task is to assign to each request a connected loop-free path of arcs from the start node to the end node. This is to be done such that the total cost for those paths is minimized, without the load on any arc exceeding its capacity, with the load defined as the number of requests sharing it. A three-request problem example is shown in Figure 1. Not all problems of this kind are solvable. A reliable algorithm should be able to recognize and signal a failure, so that proper measures can be taken.

3 The Bellman-Ford Algorithm in the Mean Field Language

Prior to dealing with the MSPP, we revisit the simpler SPP and demonstrate how the BF algorithm can be recast in a Potts MF language. This formulation will be the starting point for designing a Potts MF approach to MSPP.

In the SPP, there is only a single request, from a to b , and the capacity constraints are irrelevant. The task is simply to find the shortest path from a to b . In the BF algorithm (Bertsekas & Gallager, 1987) this is achieved by minimizing the path lengths D_i from every node i to b , by iterating

$$D_i \rightarrow \min_j (d_{ij} + D_j), \quad i \neq b \quad (3.1)$$

and keeping track of the involved arcs ij . D_b is fixed to zero by definition. The resulting solution defines a spanning tree rooted at b . In particular, D_a

is determined, and the minimal path from a to b is easily extracted from the spanning tree.

If no arc exists from node i to j , d_{ij} could formally be defined to be infinite; in practice it is more convenient to restrict j in equation 3.1 to the actual neighbors of i , reachable via an arc from i .

Equation 3.1 can be rewritten as

$$D_i = \sum_j v_{ij} E_{ij} \equiv \sum_j v_{ij} (d_{ij} + D_j), \quad (3.2)$$

in terms of a Potts spin v_i for every node $i \neq b$, with components v_{ij} taking the value 1 for the j with the smallest local energy E_{ij} , and 0 for the others (winner takes all). Note the distinct philosophy here: each node i minimizes its own local energy $D_i = \min_j E_{ij}$, rather than all nodes striving to minimize some global objective function.

An MF version of equation 3.2 is obtained by using for v_i its thermal average in the MF approximation, defined by

$$v_{ij} = \frac{e^{-E_{ij}/T}}{\sum_k e^{-E_{ik}/T}}, \quad (3.3)$$

where j and k are neighbors of i and T is an artificial temperature. Note that each Potts MF neuron v_i obeys the normalization condition

$$\sum_j v_{ij} = 1, \quad i \neq b. \quad (3.4)$$

Thus, v_{ij} can be interpreted as a probability for node i to choose j as a continuation node.

At a nonzero temperature, iteration of equations 3.2 and 3.3 can be viewed as a fuzzy implementation of the BF algorithm, while in the $T \rightarrow 0$ limit, the neurons are forced on-shell,¹ that is, $v_{ij} \rightarrow 1$ (for the minimizing j) or 0 (for the rest), and proper BF is recovered.

Given this obvious neural recast of the BF algorithm in terms of Potts neurons, it is somewhat surprising that nonexact neural approaches based on Ising spins have been advocated in the literature (Thomopolous, Zhang, & Wann, 1991).

4 The Potts Mean-Field Approach to MSPP

The Potts MF formulation of the Bellman-Ford algorithm for SPP (equations 3.2 and 3.3) is a suitable starting point for approaching MSPP. We will

¹ This notation originates from physics, where “on-shell” and “off-shell” denote whether particles are real or virtual, respectively.

stick with the philosophy inherited from BF of focusing on independent local energies, in contrast to what has become standard when using feedback neural networks for resource allocation problems. This represents a novel strategy. Thus, we introduce a separate Potts system, $\{v_{ij}^r\}$, for each request r , with basic local energies E_{ij} as before representing distances to the end node b_r . In addition, we will need energy terms E_{ij}^{load} for the load constraints, to be discussed later; this introduces an interaction among the Potts systems.

This formulation introduces the possibility of undesired loop formation, since forming a loop might induce less energy penalty than violating a load constraint. As will be discussed below, such loops can be suppressed by suitable penalty terms and by adding a possibility for each proper node to connect, by means of an artificial escape arc, to an artificial escape node, for each request connecting directly to the end node. This enables a “give-up” state for an unresolvable situation, signaled by some path containing the escape node. The cost for “giving up” must be larger than that for any legal path. Therefore, the cost of each escape arc is set to the sum of the costs of the proper arcs, while the corresponding capacity is chosen large enough to be able to host all the requests.

In order to terminate the path for a request r , its end node must be a sink for the corresponding Potts system. Consequently, there will be no Potts neuron $\mathbf{v}_{b_r}^r$ associated with it.

In order to construct appropriate penalty terms, a propagator matrix will be used. This technique has proved to be a powerful tool in neural optimization for problems with a nontrivial topology (Lagerholm, et al., 1997). In particular, it will be crucial for extracting properties of the fuzzy paths defined by the MF approach at finite T .

4.1 Path Extraction and the Propagator. The normalization condition (see equation 3.4) ensures that for each request r , precisely one continuation is chosen for each node except b_r , although for $T \neq 0$, it is fuzzily distributed over the available neighbors. On shell, the path from start to end node is trivial to extract. One follows the $v_{ij}^r = 1$ path starting from the start node. However, for $T \neq 0$ a more refined path extraction mechanism is needed. This is provided by a propagator matrix (Lagerholm et al., 1997) \mathbf{P}^r for each request r , defined by:

$$P_{ij}^r = \left[(\mathbf{1} - \mathbf{v}^r)^{-1} \right]_{ij} = \delta_{ij} + v_{ij}^r + \sum_k v_{ik}^r v_{kj}^r + \sum_{kl} v_{ik}^r v_{kl}^r v_{lj}^r + \dots \quad (4.1)$$

For a graphical representation, see Figure 2. On shell, it is easy to see that P_{ij}^r can be interpreted as the number of paths from i to j defined by the Potts



Figure 2: Graphical illustration of the expansion, equation 3.5, of the propagator \mathbf{P}^r with respect to path length. A dot represents a node i , and a line a single node-to-node step \mathbf{v}_{ij}^r .

neurons associated with r ; similarly, the elements of the matrix square $(P^r)_{ij}^2$ are related to the number of arcs used in those paths.²

In particular, we have in the absence of loops ($\Rightarrow P_{ii}^r \equiv 1$),

$$P_{a,i}^r = \begin{cases} 1, & \text{if node } i \text{ appears in the path } a_r \rightarrow b_r, \\ 0, & \text{otherwise,} \end{cases}$$

identifying the arcs used in the paths. Off shell, these interpretations are still valid in a probabilistic sense. Thus, a probabilistic measure of how much node i participates in the path $a_r \rightarrow b_r$ is given (naively) by $P_{a,i}^r$ (since $P_{i,b_r}^r \equiv 1$).³ Dividing by P_{ii}^r to correct for a possible normalization error due to a small loop contamination yields the path factor

$$F_i^r \equiv \frac{P_{a,i}^r}{P_{ii}^r} \leq 1, \quad (4.2)$$

which can be used to filter out the nodes of the fuzzy paths defined by the MF neurons.

4.2 Load and Loop Control. Armed with the propagator formalism, we proceed to set up the penalty terms, to be added to the energies E_{ij} corresponding to equation 3.2.

In order to avoid complications from self-interactions in the local energies, independently of T , we follow Ohlsson, Peterson, and Söderberg (1993), and define the penalty terms based on analyzing the result of setting one component \mathbf{v}_{ij}^r of the neuron \mathbf{v}_i^r at a time to one, with the other components set to zero, as compared to a reference state with all components set to zero.

The total load L_{ij} on an arc ij is the sum of contributions from the different requests,

$$L_{ij} = \sum_r L_{ij}^r. \quad (4.3)$$

² More precisely, the number of arcs is given by $P^2 - P$.

³ The path must eventually end in the sink b_r . *Proof:* Multiply equation 3.4 for \mathbf{v}^r by P_{ki}^r , sum over i , and use the matrix identity $\mathbf{P}^r = \mathbf{1} + \mathbf{P}^r \mathbf{v}^r$.

The contribution from request r can be expressed as

$$L_{ij}^r = F_i^r v_{ij}^r, \quad (4.4)$$

where we have used the path factor of equation 4.2.

The load constraints, $L_{ij} \leq C_{ij}$, define a set of inequality constraints. In the realm of feedback neural networks, such constraints have been successfully handled by means of step functions (Ohlsson et al., 1993). For a given request r , the overloading of the arc ij due to the other requests is given by

$$O(X) \equiv X \Theta(X), \quad (4.5)$$

where $\Theta()$ is the Heaviside step function, and X is given by

$$X \equiv L_{ij} - L_{ij}^r - C_{ij}. \quad (4.6)$$

If the arc were also to be used by r , the overloading would increase to $O(X+1)$, and a suitable overloading penalty can be defined as the difference,

$$E_{ij}^{\text{load}} = O(X+1) - O(X). \quad (4.7)$$

The number of loops introduced by connecting $i \rightarrow j$ can be expressed as

$$Y \equiv \frac{P_{ji}^r}{P_{ii}^r} \leq 1, \quad (4.8)$$

and we choose as a loop-suppression term,

$$E_{ij}^{\text{loop}} = \frac{Y}{1-Y}. \quad (4.9)$$

The generalization of the local energy in equation 3.2 to the multiple request case now reads, for a particular request r ,

$$E_{ij} = d_{ij} + D_j^r + \alpha E_{ij}^{\text{load}} + \gamma E_{ij}^{\text{loop}}, \quad (4.10)$$

with the added terms based on equations 4.7 and 4.9. The resulting algorithm allows for a wide range of choices of the coefficients α and γ without severely changing the performance.

4.3 Updating Equations. All neurons are repeatedly updated, with a slow annealing in T . For each request r and each node i , the corresponding neuron v_i^r is updated according to

$$v_{ij}^r = \frac{e^{-E_{ij}/T}}{\sum_k e^{-E_{ik}/T}}, \quad (4.11)$$

with E_{ij} given by equation 4.10. The corresponding cost D_i^r from node i to the end node is then updated, in the BF spirit, as

$$D_i^r \rightarrow \sum_j v_{ij}^r E_{ij}. \quad (4.12)$$

In principle, the corresponding update of the propagator could be done using an exact incremental matrix inversion scheme like Sherman-Morrison (Press, Flannery, Teukolsky, & Vetterling, 1986). We prefer, though, to let local changes propagate through the network, in analogy to the update of D_i^r . Thus, only the i th row of \mathbf{P}^r is updated:

$$P_{im}^r \rightarrow \delta_{im} + \sum_j v_{ij}^r P_{jm}^r, \text{ for all } m. \quad (4.13)$$

This gives a convergence toward the exact inverse, which turns out to be good enough. The advantage of this method is twofold: it is faster, and all information needed is local to the relevant node i and its neighbors j (assuming each node to keep track of its own row of \mathbf{P}^r). Details of the algorithmic steps and the initialization can be found in the appendix.

4.4 Test Problems and Explorations. In order to test the Potts MF method, we have generated a set of challenging synthetic problems. The most important parameters governing the difficulty of a problem, apart from network size and connectivity, are the number of requests and the average arc capacity. In cases where all the arcs are able to host all the requests, $C_{ij} \geq N_R$, the problem is separable into an independent SPP for each request, which can be solved using the BF algorithm. We have therefore chosen to work with rather tight arc capacities.

For each problem, we generate a random connected network, where every node has at least one path to all other nodes. To that end, all nodes are first connected in a random spanning tree. Additional links (creating loops) are then randomly placed. Every arc is given a random integer capacity between 1 and 6 and a random cost in the interval $(0, 1)$. The desired number of random requests is generated, in terms of start and end nodes. An example of a generated test problem is shown in Figure 3.

This procedure does not automatically yield a solvable problem, where all requests can be fulfilled simultaneously without violating any constraint. In principle, solvability could be built into the problem generator, but here we adopt another strategy. We attempt to solve small problems (see Table 1) exactly with a BB algorithm; those not solved within a certain amount of CPU time are disregarded. In principle, this method could introduce a bias toward simple problems. However, for the smaller problem sizes considered, only a tiny fraction of the problem candidates are “timed-out,” except for those defining the last row in Table 1, where about one-third are disre-

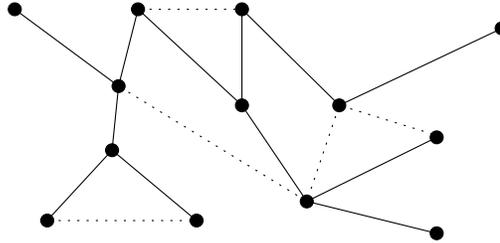


Figure 3: An example network with 13 nodes and 17 links. The solid lines define the initial spanning tree, and the dotted lines represent the additional links.

Table 1: Results for Small Problems.

Nodes	Links	Requests	$\langle S \rangle$	Legal_{MF}	$\langle \Delta_{\text{BB}} \rangle$	$\langle \text{CPU time} \rangle$
5	10	5	14	100.0%	0.003	0.1
5	10	10	28	99.9	0.002	0.2
10	15	10	28	100.0	0.004	0.4
10	20	10	48	100.0	0.003	0.5
15	20	10	27	99.8	0.03	0.5
15	20	15	40	99.9	0.06	0.7

Note: 1000 problems of each size are probed. Only legal entries are used when calculating the averages $\langle \Delta_{\text{BB}} \rangle$ and $\langle \text{CPU time} \rangle$. $\langle \text{CPU time} \rangle$ refers to the MF Potts approach using a DEC Alpha 2000, and is given in seconds. Typical times for the BB method are around 600 seconds.

garded. For larger problems (see Table 2), BB cannot be used, and we do not know in general whether they are solvable.

Table 2: Results for Large Problems.

N	N_L	N_R	Reduction	$\text{legal}_{\text{IBF}}$	$\langle \Delta_{\text{IBF}} \rangle$	$\text{legal}_{\text{SBF}}$	$\langle \Delta_{\text{SBF}} \rangle$	legal_{MF}
50	200	20	1.00	4%	0.000	100%	-0.019	100%
80	200	20	0.97	8	0.000	98	-0.021	98
100	200	20	0.91	7	0.000	96	-0.015	96
100	200	50	0.90	0	---	82	-0.043	82
100	200	100	0.91	0	---	22	-0.092	30
100	300	100	0.99	0	---	83	-0.066	84
100	400	100	1.00	0	---	99	-0.063	100

Note: 100 problems of each kind are probed. The averages, $\langle \Delta_X \rangle$, are for legal entries. The reduction is defined as $N_R N N_L$ after the decomposition divided by the value before the decomposition. Typical CPU times for MF are $O(10)$ minutes on a DEC Alpha 2000.

Prior to attacking a problem, a decomposition into independent subproblems is attempted, to reduce complexity. The decomposition uses the fact that a network typically consists of subnetworks, connected to each other by single nodes. The complete problem can then be solved by independently solving the implied subproblem for each subnetwork.

The CPU demand of the nonreduced MF approach scales like $N_R N N_L$, while the required computer time for a decomposed problem is dominated by the largest subproblem. In Table 2 we indicate the effect on the CPU demand of the decomposition for the large test problems.

The performance of the MF Potts approach is probed by measuring the relative excess path length as compared to a reference algorithm (X),

$$\Delta_X = \frac{D_{MF} - D_X}{D_X}, \quad (4.14)$$

where D_{MF} and D_X are the total path costs resulting from MF and the reference algorithm, respectively.

4.5 Comparison with an Exact Method. For the smaller problems, the BB algorithm is used as a reference algorithm. The results are displayed in Table 1, together with characteristics of the generated test problems. As a measure of the complexity of a problem, we use the entropy, S , defined as the logarithm of the total number of possible configurations, disregarding load constraints. Table 1 indicates excellent performance of the MF Potts approach, with respect to giving rise to legal solutions with good quality, with a very modest computational demand.

4.6 Comparing with Other Approximative Methods. The computer demand for the BB algorithm grows rapidly. For problems larger than those of Table 1, we have chosen as reference algorithms two BF-inspired heuristics:

- Independent Bellman-Ford (IBF). Each request is taken to define an independent SPP, which is solved using BF. Thus, the load constraints are neglected, and the resulting solution is not necessarily legal. If legal, it gives the global energy minimum.
- Sequential Bellman-Ford (SBF). The requests are served in a sequential fashion using the BF algorithm. When the maximum capacity of an arc is reached, its use is prohibited for the remaining requests. It does not always find a solution, even to a solvable problem, and when it does, it is not necessarily the global minimum.

The early ARPANET (Bertsekas & Gallager, 1987) (Internet) protocol was based on an algorithm similar to IBF (but with a dynamically updated arc cost based on traffic congestion). The results are displayed in Table 2. In the

few cases where the global minimum is known (from IBF), it is found by the MF algorithm. In the other cases, we conclude from Table 2 that MF on the average performs better than SBF. When the MF cannot find a solution, this is signaled by requests routed via the escape node.⁴ In all such cases, also IBF and SBF have failed.

5 Summary and Outlook

We have developed a Potts MF neural network algorithm for finding approximate solutions to the MSPP. The starting point is a Potts MF recast of the exact Bellman-Ford algorithm for the simpler single SPP. This approach is then extended to the MSPP by using several Potts networks, one for each request.

Complications of topological nature are successfully handled by means of a convenient propagator approach, which is crucial for the following issues:

- The MF approach yields at $T \neq 0$ fuzzy spanning trees, from which the propagator is used to extract the loads corresponding to the fuzzy paths, needed for the interaction between different requests.
- Loops are suppressed by energy terms, based on the propagator.

To open up escape paths for unresolvable situations, an auxiliary arc to an escape node is introduced for each proper node. The method is local in that only information available from neighboring nodes is required for the updates. This attractive feature, inherited from the Bellman-Ford algorithm, facilitates a distributed implementation.

The computational demand of the method is modest. The CPU time scales as $N_R N N_L$. With fixed connectivity, this corresponds to $N_R N^2$, whereas for the worst case of full connectivity, it yields $N_R N^3$.

The performance of the algorithm is tested on a set of synthetic challenging problems. This is done by comparing to exact results from a BB method for smaller problems and comparing to results from independent and sequential Bellman-Ford approaches for larger problems. The comparisons show that the Potts MF approach with very few exceptions yields very good approximate solutions at modest computer time consumption. The method is now being generalized to other routing problems (Häkkinen, Lagerholm, Peterson, & Söderberg, 1998).

⁴ This could be used in dynamical cases as a decision tool for which requests should be put on hold.

Appendix: The Potts MF Algorithm

A.1 Initialization. The initial temperature T_0 is first set to $T_0 = 50$. If the saturation Σ ,

$$\Sigma \equiv \frac{1}{N_R(N-1)} \sum_{i \neq b_r} \mathbf{v}_i^2, \quad (\text{A.1})$$

has changed more than 10% after all neurons have been updated once, then the system is reinitialized, with $T_0 \rightarrow 2T_0$.

For all nodes (except the end and escape nodes), the corresponding Potts neurons are initialized in accordance with the high temperature limit, that is,

$$v_{i,j}^r = 1/n_i, \quad (\text{A.2})$$

for all n_i neighbors j (including the escape node) of i . P_{ij}^r and D_i^r are initialized consistently with equation A.2.

A.2 Iteration. Until $T \leq T_f$ or $\Sigma \geq \Sigma_f$ do:

- For every request r do:
 1. For every node i except b_r and the escape node:
 - (a) Update \mathbf{v}_i^r (see equations 3.3 and 4.10).
 - (b) Update D_i^r (see equation 4.12).
 - (c) Update \mathbf{P}_i^r (see equation 4.13).
 2. Update L_{ij} .
- Decrease the temperature: $T = kT$.

We have used the parameter values $k = 0.9$, $T_f = 0.0001$, and $\Sigma_f = 0.99999$. For the energy coefficients in equation 4.10, we have consistently used $\alpha = 1$ and $\gamma = 5$.

References

- Bellman, R. (1958). On a routing problem. *Quarterly of Applied Mathematics*, 16, 87.
- Bertsekas, D., & Gallager, R. (1987). *Data networks*. Englewood Cliffs, NJ: Prentice-Hall.
- Boyan, J. A., & Littman, M. L. (1994). Packet routing in dynamically changing networks: A reinforcement learning approach. In J. D. Cowan, G. Teauso, & J. Alspector (Eds.), *Advances in neural information processing systems*, 6. San Mateo, CA: Morgan Kaufman.

- Gislén, L., Peterson, C., & Söderberg, B. (1992). Complex scheduling with Potts neural networks. *Neural Computation*, 4, 805.
- Häkkinen, J., Lagerholm, M., Peterson, C., & Söderberg, B. (1998). In preparation.
- Lagerholm, M., Peterson, C., & Söderberg, B. (1997). Airline crew scheduling with potts neurons. *Neural Computation*, 9, 1627.
- Ohlsson, M., Peterson, C., & Söderberg, B. (1993). Neural networks for optimization problems with inequality constraints—the knapsack problem. *Neural Computation*, 5, 331.
- Peterson, C., & Söderberg, B. (1989). A new method for mapping optimization problems onto neural networks. *International Journal of Neural Systems*, 1, 3.
- Press, W. P., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T. (1986). *Numerical recipes: The art of scientific computing*. Cambridge: Cambridge University Press.
- Thomopolous, S. C. A., Zhang, L., & Wann, C. D. (1991). Neural network implementation of the shortest path algorithm for traffic routing in communication networks. *Proceedings of the IEEE International Joint Conference on Neural Networks*. Singapore.

Received March 14, 1997; accepted December 10, 1997.