c3

The Principle of Least Cost, The Friction of Conflict and The Quantum Mechanics of Command and Control

B D Bramson & T R Field
DERA, St Andrews Road, Malvern, Worcs, WR14 3PS

22 August 1997

Contents

1 Setting the scene

It is the year 2020. The countries of East and West Phalia have summoned assistance from the UN to coordinate a joint mission to deal with the poisonous gas cloud that has appeared over their common border. Each country will launch a swarm of "nano-beasties", miniature airborne platforms bristling with sensors, computing power and laboratories on chips, equipped with receivers tuned into the GPS (Global Positioning System) and the GBS (Global Broadcast System), and propelled by on-board, power plants.

And so the time has come and the swarms are launched.

Their purpose: to rendezvous at 0300h on 1 Apr 2020 as close to the cloud centre as possible, thence to analyse and destroy it.

Problem 1: the fuel requirements for the mission are severe.

Problem 2: North Phalia has not agreed to the use of its airspace which must therefore be regarded as a no-go area; beasties that stray into no-go areas run the risk of destruction.

Problem 3: any swarm of beasties will be subject to continuous, random disturbances owing to the unpredictability of local airflow.

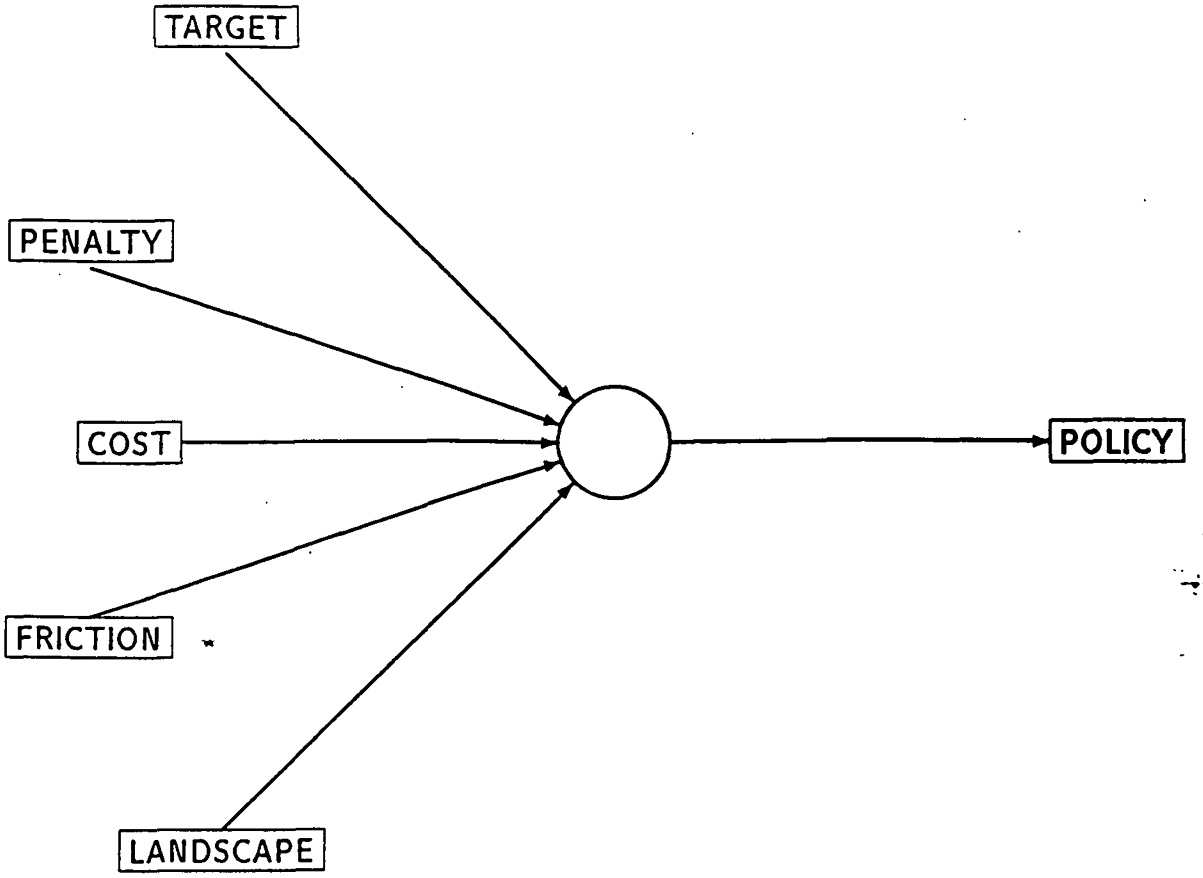

Figure 1: A policy is a velocity field $\mathbf{v}(t, \mathbf{x})$ on spacetime: it allows for random disturbances and cannot be portrayed as a schedule.

The UN has access to all intelligence concerning the location of the gas cloud and the locations, geometry, and severity of all no-go areas. It is clear, however, particularly in view of Problem 3, that specifying a schedule, an *a priori* sequence of instructions, to be followed by swarm members will be useless.

2 Planning a rendezvous

If schedules are out, what is the alternative? Can we define an approach that will provide the swarm members with the "best" route to the rendezvous point? Clearly, we need to minimise both the fuel used and the likely level of attrition, two factors that may pull in opposite directions. So, let us conceive of a planning process (figure 1) that inputs the following:

1. the target location $\mathbf{X}$ and rendezvous time T;
2. the penalty $P(\mathbf{x}, \mathbf{X})$ to be paid at time T, given the extent to which the target is missed;
3. details of the military landscape, involving, for example, the nature of terrain and military force concentrations, encoded as a potential function $V(\mathbf{x})$;
4. an expression $L(\mathbf{x}, \mathbf{v})$ for the local rate of expenditure involving the velocity of motion of the platforms and their locations;
5. a constant parameter σ that accounts for the expected degree of random disturbance.

The process shall output not a schedule but rather a policy to be broadcast and updated as necessary. The policy says to each platform "If you find yourself at location $\mathbf{x}$ at time t, adopt the motion determined by the fluid flow whose velocity is given by $\mathbf{v}(t,\mathbf{x})$". Adopting a policy rather than a schedule allows implicitly for the effects of random disturbances.

This is very much the approach of Rosenbrock [1] who layed the foundation for the line of research that we shall outline next.

3 The approach

Let us consider a beastie whose location at time t is described by a probability density $\rho(t,\mathbf{x})$ and whose motion is specified by a velocity field $\mathbf{v}(t,\mathbf{x})$ that determines the policy. For the moment ignore attrition. To account for friction, superimpose on the motion a random walk that is isotropic and homogeneous. (In one spatial dimension, toss a coin every second and if it's heads walk north a metre and if it's tails walk south.) The classical equation of continuity involving $\rho(t,\mathbf{x})$ acquires an extra term involving the Laplacian of the density: equivalently this is the diffusion equation with variable drift. (See also references to the heat equation, Fokker-Planck equation and Brownian motion [2, 3].)

So far what we have said applies to any policy determined by the flow field $\mathbf{v}(t,\mathbf{x})$. Now we are all accountants: so we had better raise the issue of cost. We shall assume that as the beastie follows a given path (*modulo* its random walking) it incurs a travel cost. The rate of cost expenditure is assumed to take the form $L(\mathbf{x},\mathbf{v}) = \frac{1}{2}m\mathbf{v}^2 - V(\mathbf{x})$, for some constant m, and this expression is integrated along the possible trajectories of motion. (The potential term V, which will usually be negative, is intended to encode terrain, culture and military force concentrations.) At time T the beastie will incur an additional cost, a penalty $P(\mathbf{x},\mathbf{X})$, depending on its miss distance from the target $\mathbf{X}$. At any point $(t,\mathbf{x})$ in spacetime we may associate a remaining cost for the mission $S(t,\mathbf{x})$. This is the cost that will be expended from time t to time T if the beastie follows the given policy from $\mathbf{x}$. Of course the rate of change of $S(t,\mathbf{x})$ along the flow lines given by the policy is intimately related to $L(\mathbf{x},\mathbf{v})$ but there is an extra dispersive term that represents friction. Indeed, this dispersion is an entropy increasing process while the command and control, typified by the policy $\mathbf{v}(t,\mathbf{x})$, is an entropy reducing process. When $t = T$, the mission is over and we must identify $S(T,\mathbf{x})$ with the penalty $P(\mathbf{x},\mathbf{X})$. This is clearly a final value problem, one posed in terms of Rosenbrock's purposive myth rather than the initial value problem that typifies the causal myth.

Now, when we apply the principle of least cost an extraordinary state of affairs emerges [4]. The beastie's probability density decomposes into the product of a pair of real "wave functions". Each satisfies a linear, partial differential equation. One of these propagates backwards in time from the data supplied by the penalty function at time T, while its companion propagates forwards from data defined by the beastie's initial state. The entire structure is analogous to what one finds in quantum mechanics. Indeed the equations of propagation are real versions of Schrödinger's equation. However, all this disappears the moment we switch off the random disturbances.

4 Outline of mathematical details

The probability density of a given beastie satisfies the *Fokker-Planck equation* [2]

$$\partial\rho/\partial t + \nabla.(\rho\mathbf{v}) = \sigma\nabla^2\rho. \qquad (4.1)$$

The vector field $\mathbf{v}(t,\mathbf{x})$ here defines the *mean forward drift velocity* [3] and it reveals the platform's most likely position a short time into the future given knowledge of its current position. A corresponding velocity $\bar{\mathbf{v}}(t,\mathbf{x})$, the *mean backward drift*, which satisfies

$$\partial\rho/\partial t + \nabla.(\rho\bar{\mathbf{v}}) = -\sigma\nabla^2\rho, \tag{4.2}$$

reveals the platform's most likely position a short time into the past given its current position. The velocity used by most fluid dynamicists is actually the average of these two. We claim, however, following Rosenbrock, that a *policy* should deal with the mean *forward* drift.

At time t, the expected cost remaining for the mission is assumed to take the form

$$\int d^3\mathbf{x}\,\rho(t,\mathbf{x})S(t,\mathbf{x}) = \int_t^T ds \int d^3\mathbf{x}\,\rho(s,\mathbf{x})L(\mathbf{x},\mathbf{v}(s,\mathbf{x})) + \int d^3\mathbf{x}\,\rho(T,\mathbf{x})P(\mathbf{x},\mathbf{X}), \tag{4.3}$$

whatever the choice of ρ. The equation of motion for S, posed as a final value problem, is

$$\left.\begin{aligned} \partial S/\partial t + \mathbf{v}.\nabla S + \sigma\nabla^2 S &= -L, \\ S(T,\mathbf{x}) &= P(\mathbf{x},\mathbf{X}). \end{aligned}\right\} \tag{4.4}$$

If we choose L to take the form $\frac{1}{2}m\mathbf{v}^2 - V(\mathbf{x})$, it follows that the policy minimising the cost is characterised by the relationship

$$\nabla S + m\mathbf{v} = 0. \tag{4.5}$$

Defining a real wave function ψ via the coordinate transformation

$$S = -h\log\psi, \qquad [h = 2m\sigma] \tag{4.6}$$

then yields the time-reversed diffusion equation for ψ:

$$\left.\begin{aligned} -h\partial\psi/\partial t &= (h^2/2m)\nabla^2\psi + V\psi, \\ \psi(T,\mathbf{x}) &= \exp(-P(\mathbf{x},\mathbf{X})/h). \end{aligned}\right\} \tag{4.7}$$

We define a second wave function $\bar{\psi}$ by writing $\rho = \bar{\psi}\psi$. Then demanding the equation of continuity for ρ yields the diffusion equation for $\bar{\psi}$:

$$\left.\begin{aligned} h\partial\bar{\psi}/\partial t &= (h^2/2m)\nabla^2\bar{\psi} + V\bar{\psi}, \\ \bar{\psi}(0,\mathbf{x}) &= \rho(0,\mathbf{x})/\psi(0,\mathbf{x}). \end{aligned}\right\} \tag{4.8}$$

The similarity of the equations (4.7) and (4.8) to those found in quantum mechanics is striking. Indeed, a complex structure (cf [5]) can be shown to exist notwithstanding the fact that the wave functions are real. However, the two formalisms are not identical. Formally, we can obtain quantum mechanics by working with a superposition $\mathbf{w}$ of $\mathbf{v}$ and $\bar{\mathbf{v}}$:

$$\mathbf{w} = z\mathbf{v} + \bar{z}\bar{\mathbf{v}}, \qquad [z = \tfrac{1}{2} - \tfrac{1}{4}i\hbar/m\sigma]. \tag{4.9}$$

In terms of $\mathbf{w}$, the equation of continuity for ρ becomes

$$\partial\rho/\partial t + \nabla.(\rho\mathbf{w}) = (-\tfrac{1}{2}i\hbar/m)\nabla^2\rho. \tag{4.10}$$

Effectively, this amounts to replacing h by $-i\hbar$ and indeed, choosing $\frac{1}{2}m\mathbf{w}^2 - V(\mathbf{x})$ instead of $\frac{1}{2}m\mathbf{v}^2 - V(\mathbf{x})$ as the rate of cost expenditure L, yields Schrödinger's equation for a single complex-valued wave function.

The difference in the two formalism amounts to the choice of the function L. This is clearly related to the way we measure the rate of cost expenditure but it may also be linked to issues of empowerment, autonomy and the collapse of the wave function. These are very deep matters which are currently undergoing further investigation.

5 Software implementation

The theory just described has been implemented in software, yielding fast, frugal simulations that have stimulated our intuition concerning the nature of fog, friction and ISTAR effectiveness. The program, written in ALGOL 68 RS and mounted on a VAX VMS platform, takes a simple ascii input file and produces LaTeX output that in turn is used to generate a graphical display. The main algorithm seeks to solve a parabolic, partial differential equation by approximating it to a finite grid and using a Taylor-series approach.

The graphical examples that follow relate to a region of terrain on which may be found the target • and a pair of no-go regions having circular structures and implemented as Gaussian distributions. Figures 2 and 3 depict the policies appropriate to zero and perfect intelligence respectively while figures 4 and 5 illustrate typical implementations of those policies.

References

[1] H H Rosenbrock. *Machines with a Purpose.* Oxford University Press, 1990.

[2] C W Gardiner. *Handbook of Stochastic Methods.* Springer-Verlag, Berlin, 1983.

[3] E Nelson. *Dynamical Theories of Brownian Motion.* Princeton University Press, 1967.

[4] B D Bramson. What mathematics for future systems? In A E R Woodcock and D F Davis, editors, *Analytic Approaches to the Study of Future Conflict*, Pearson Peacekeeping Centre, Cornwallis Park, Nova Scotia, 1996. Canadian Peacekeeping Press.

[5] T R Field. *The Quantum Complex Structure.* D Phil Thesis, Oxford University, 1996.

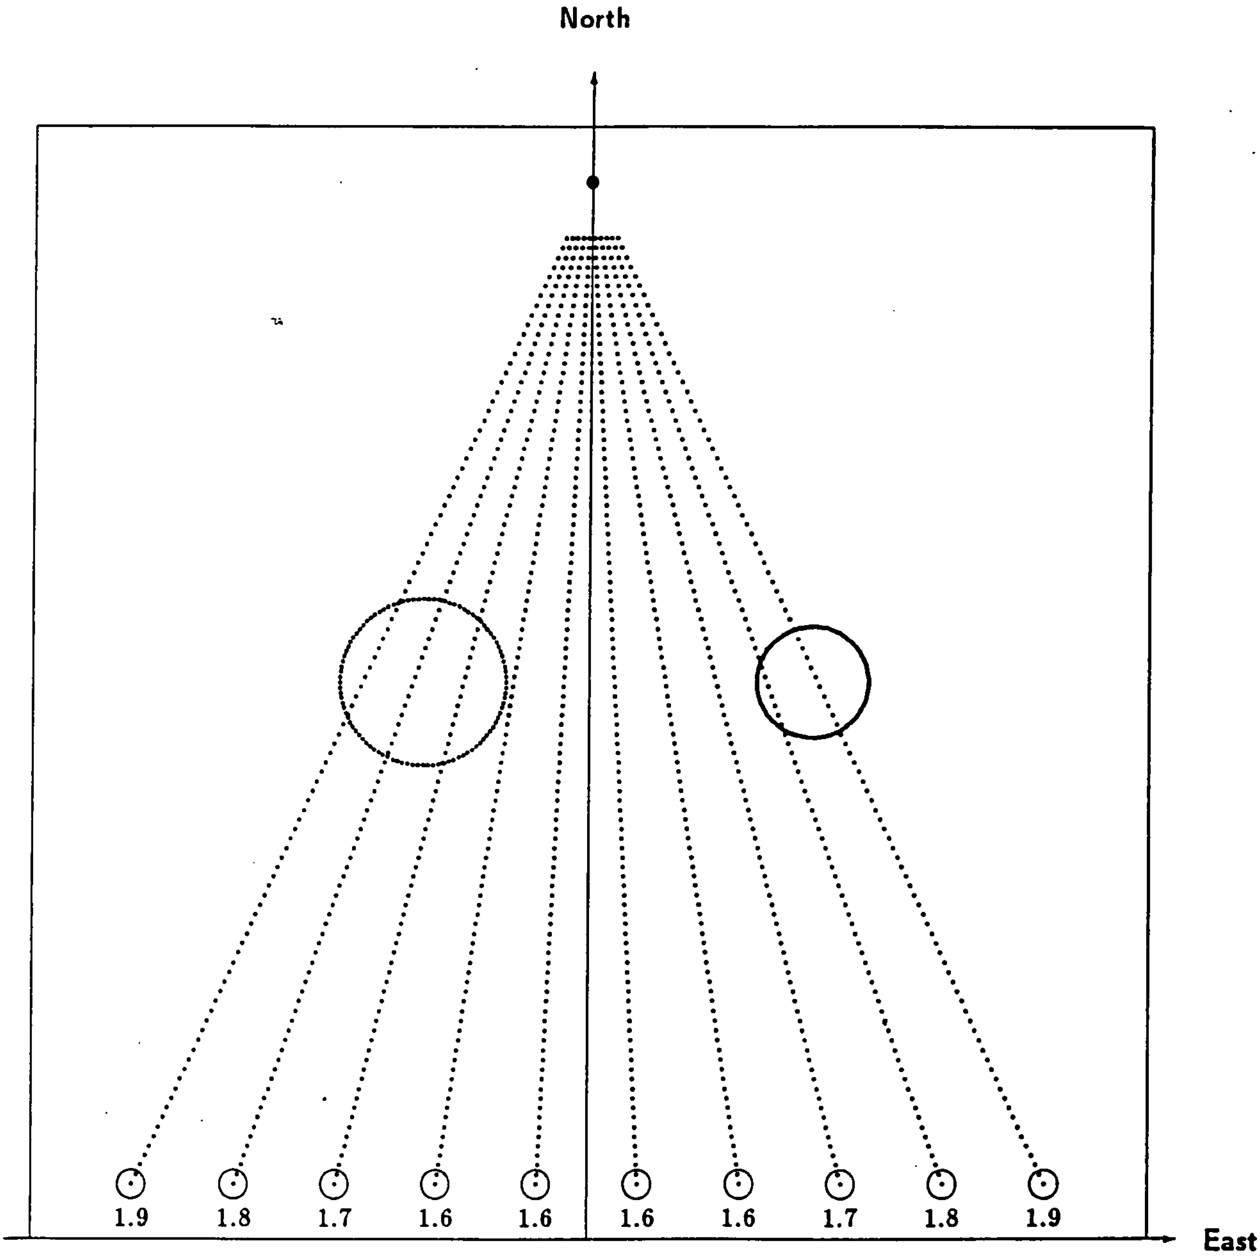

Figure 2: A policy based on zero intelligence: the estimated costs, ranging from 1.6 to 1.9, ignore the presence of danger.

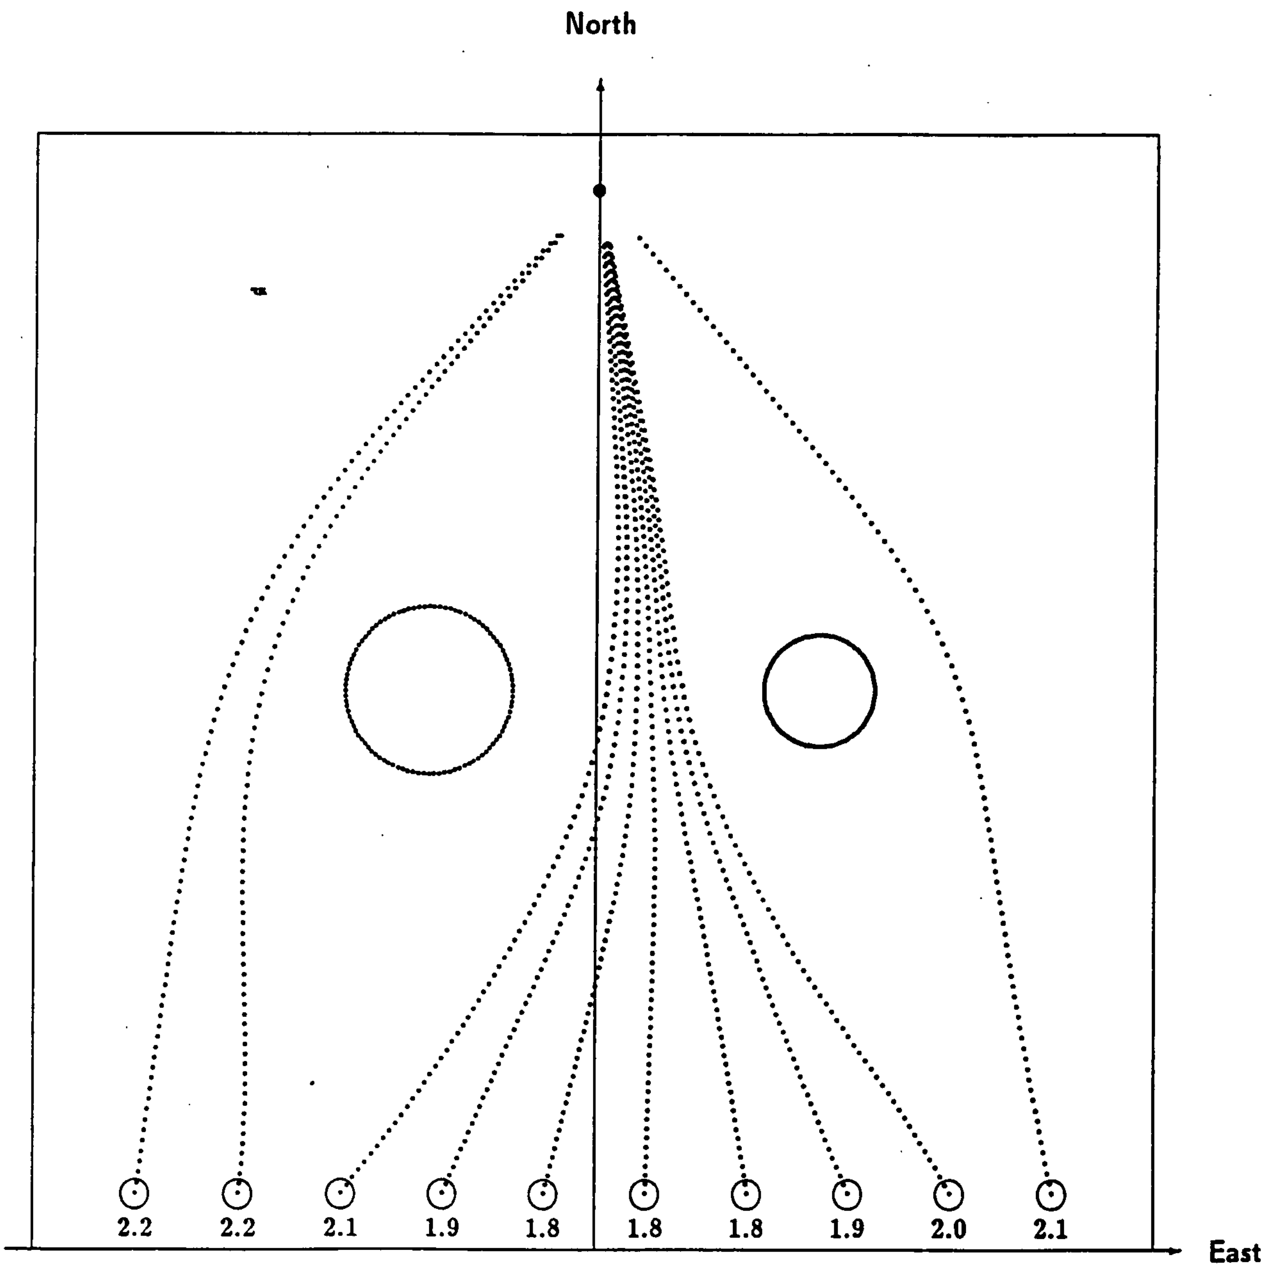

Figure 3: A policy based on perfect intelligence.

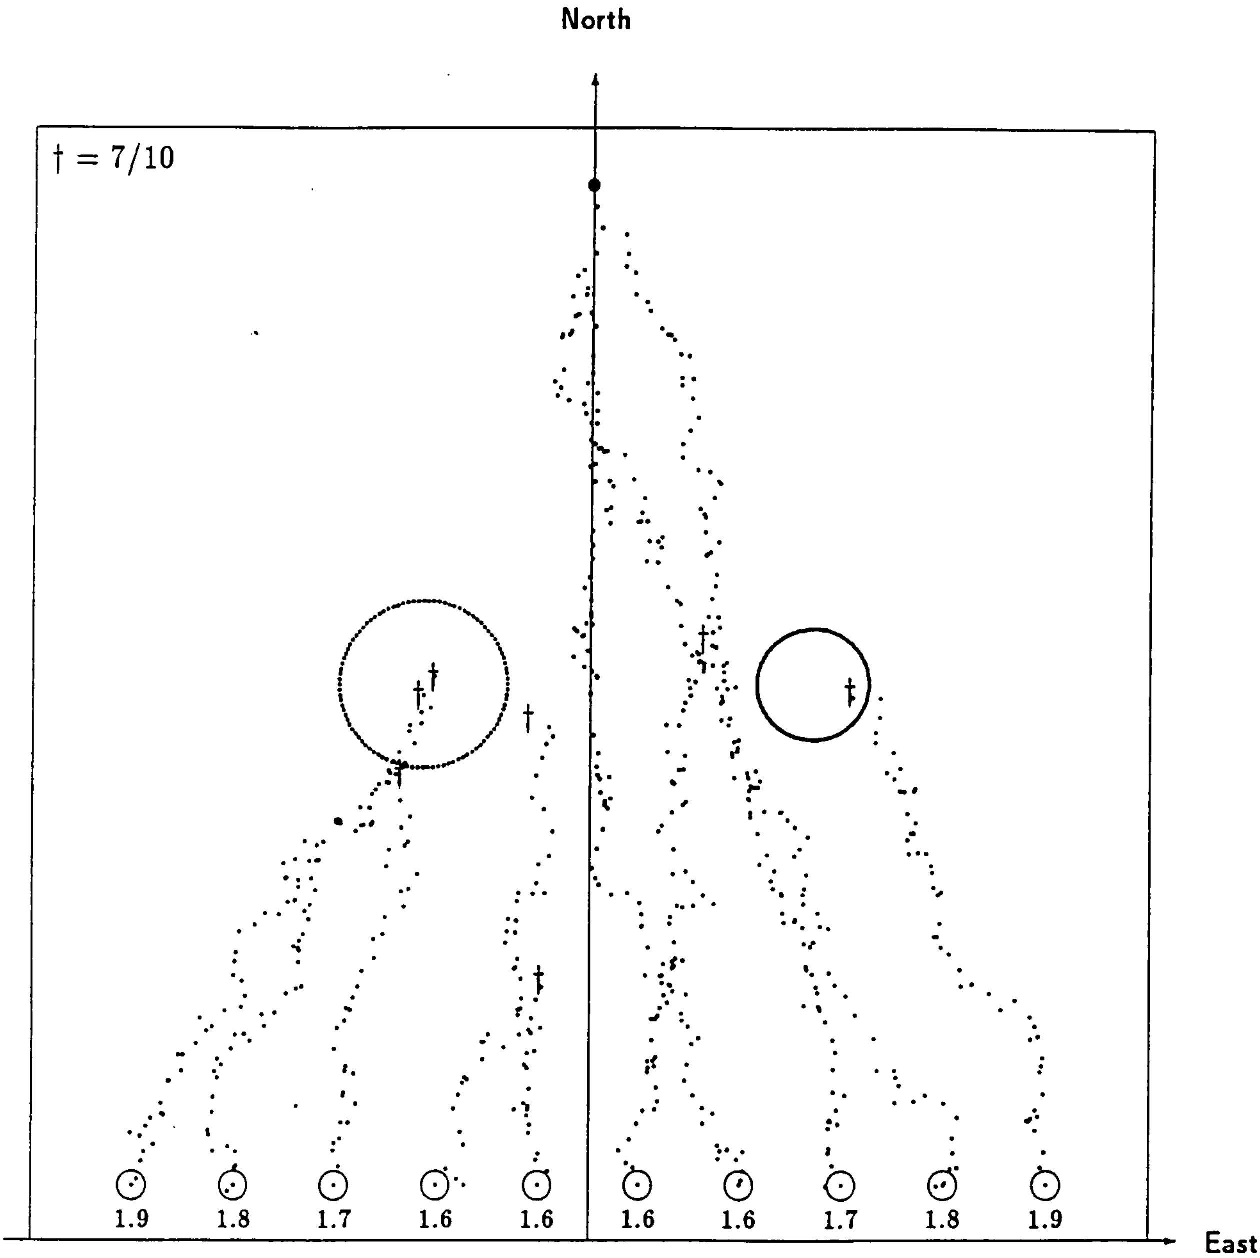

Figure 4: A typical implementation of the zero intelligence policy: 7 out of 10 platforms, indicated by †, have been destroyed.

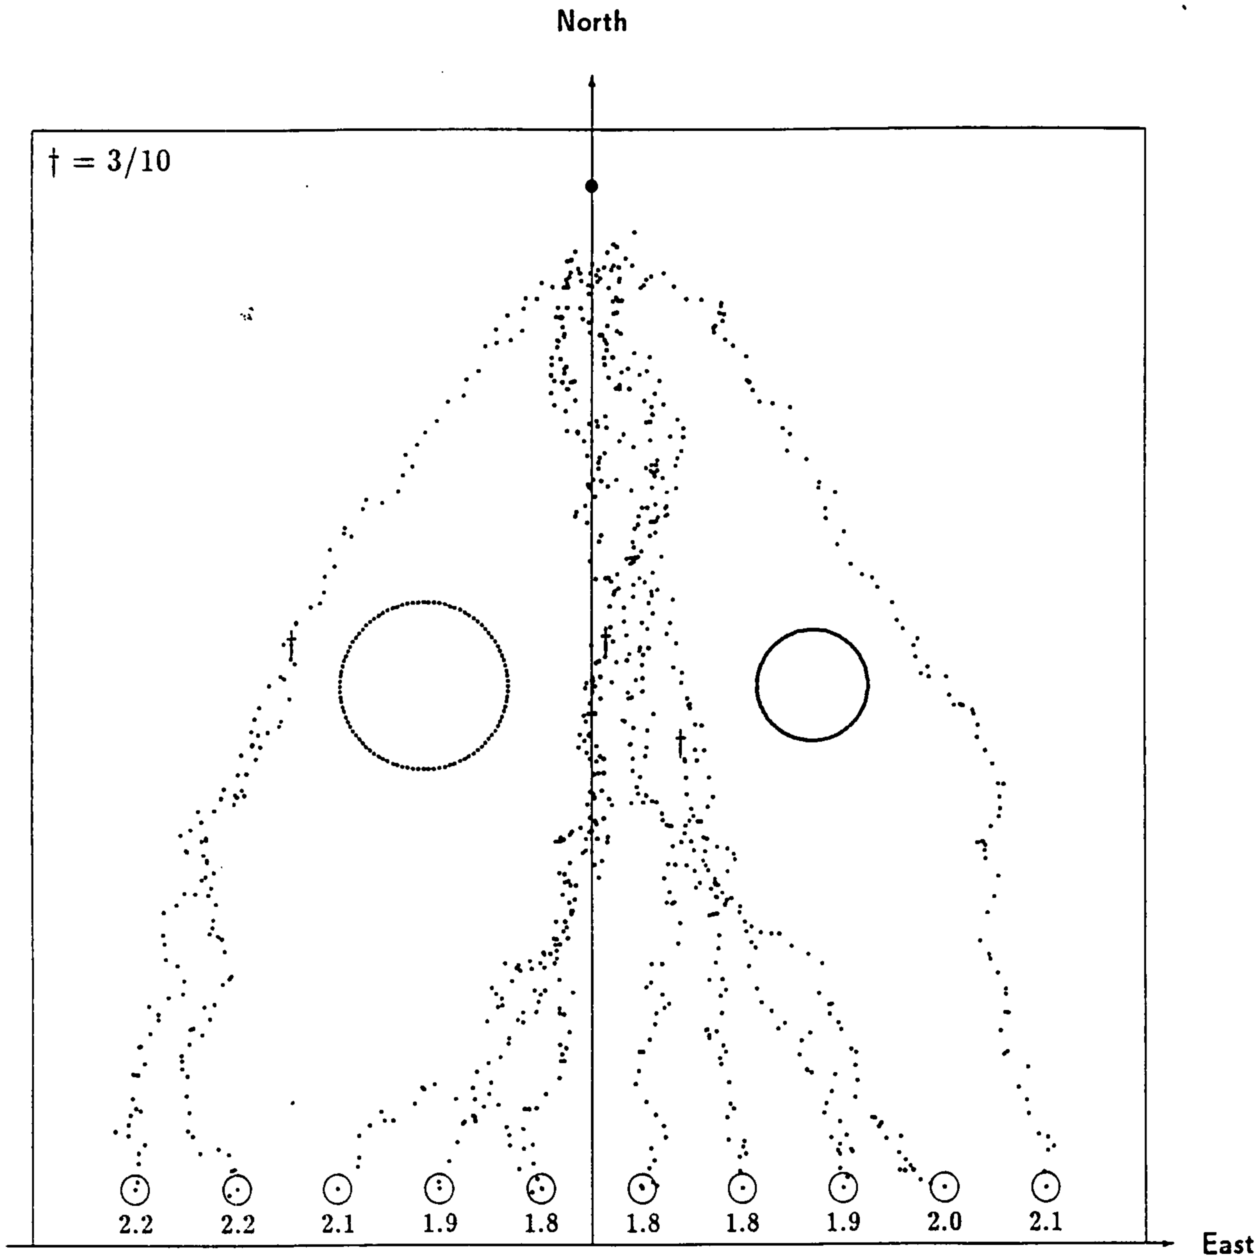

Figure 5: A typical implementation of the perfect intelligence policy: 3 out of 10 platforms, indicated by †, have been destroyed.