Smart Assignment Across Simultaneous EOD Incidents: Choosing the Right Cordon for the Right Responder

Bomb Doctor
The Long Walk Technical Series, Paper 3

June 21, 2026

Abstract

When more than one explosive hazard is live at the same time, a wrist-worn approach display can show only one incident, and something must decide which. The obvious rule, show the nearest, is often wrong, because a larger charge farther away can place a responder in a higher overpressure and fragment hazard band than a small charge underfoot. This paper develops and documents the assignment model behind The Long Walk system. It frames the problem as a spatial partition of responder positions, defines two single-responder strategies, nearest and highest-exposure, and two team strategies, independent assignment and a distinct one-responder-per-incident matching, and derives them on top of the live exposure engine described in the companion papers so that the choice the phone makes and the readout the wrist shows come from the same physics. We interpret the highest-exposure partition through weighted Voronoi geometry, which explains analytically why its boundary bends toward a small charge and away from a large one. Decision maps generated directly from the shipped code, for two and three simultaneous incidents, confirm the geometry, and a worked numerical example shows the two single-responder strategies disagreeing at specific positions. We state the limits of the model honestly: severity is a banded integer, the distinct team matching is greedy rather than globally optimal, and the model decides what a display shows, not whether a responder should move. The model is a documented engineering tool and is not operational guidance.

Model status and limitations

The quantitative models in this paper are open-literature engineering correlations implemented in software. Fragment-distance and injury correlations are *illustrative* estimates for cordon sizing, not guarantees; real outcomes depend on casing, terrain, geometry, and orientation not known in the field. All distances default to the greater of the computed value and the published doctrinal floor. Nothing here is operational guidance or a substitute for command explosive-safety authority.

1 Introduction

The first two papers in this series treat a single incident. Paper 1 sizes its cordon from charge size, scaled-distance criteria, and fragment hazard range (U.S. Department of Defense, 2008, 2019), and validates it against the bare K-factor calculator and the DHS/ATF standoff card (U.S. Department of Homeland Security & Bureau of Alcohol, Tobacco, Firearms and Explosives, 2016). Paper 2 reports a responder's live exposure to that incident as they approach, combining overpressure, fragment-hit probability, and an impulse-aware lethality estimate (Bowen et al., 1968; Kingery & Bulmash, 1984). Both assume the responder is attending to one hazard.

Real events are not always single. A vehicle-borne device may be accompanied by a secondary. A cleared lane may contain two suspect items. A training range may run several simultaneous problems. In each case a responder carries one display, that display can present one incident at a time, and a rule must select which incident to present at the responder's current position. This selection is the subject of the present paper.

1.1 Why the nearest rule fails

The simplest selection rule is to show the nearest incident. It is appealing because distance is cheap to compute and intuitive to reason about, and when all charges are comparable in size it gives the right answer. It fails precisely when charges differ, because proximity is not the same as danger. Consider a small charge ten metres away and a large charge a hundred metres away. The small charge is nearer, yet the large charge can place the responder in a higher overpressure band and a higher fragment-hit probability, as Papers 1 and 2 quantify. A rule that selects on distance alone will confidently direct the responder's attention to the lesser hazard, and it will do so most often in exactly the mixed-charge situations where the stakes of attending to the wrong hazard are highest.

1.2 Contribution and roadmap

This paper makes three contributions. First, it frames simultaneous-incident selection as a spatial partition problem and connects it to the classical theory of Voronoi diagrams and their weighted generalisations, which provides an analytic account of the partition boundaries the system produces. Second, it defines exposure-aware selection rules, for a single responder and for a team, that are computed on top of the exposure engine of Paper 2, so that selection and readout share one model. Third, it validates the rules with decision maps and a worked example generated from the shipped implementation, and it states the limits of the approach. Section 2 reviews the relevant background. Section 3 gives the formal model. Section 4 presents results. Section 5 describes the implementation and its traceability. Sections 6 and 7 give limitations and future work, and Section 8 concludes the three-paper series.

2 Background and Related Work

2.1 Assignment and facility-location problems

Deciding which of several sites a query point should be associated with is a classical question in computational geometry and operations research. In the facility-location family, a set of facilities is given and each demand point is served by one facility according to a cost, most commonly distance. The assignment problem proper seeks a matching between two sets, for example agents and tasks, that optimises a total cost. The simultaneous-incident selection problem sits between these: a single responder is a demand point to be associated with one incident, while a team of responders to be distributed across incidents is a matching problem. We treat both cases below, and we are explicit about which is solved optimally and which is solved greedily.

2.2 Voronoi partitions as the geometric lens

The nearest rule partitions the plane into Voronoi cells. Each incident owns the set of positions closer to it than to any other incident, and the boundary between two cells is the perpendicular bisector of the segment joining the two incidents. This is the standard nearest-neighbour partition of computational geometry, and it accounts for the straight midline boundary that a nearest rule produces between two equal competitors.

When the competitors are not equal, the relevant construction is a weighted Voronoi diagram. In a multiplicatively weighted diagram each site carries a weight and the association cost is distance divided by weight, so a stronger site claims territory beyond the midline and the bisecting boundary becomes a curve rather than a straight line. In a power (additively weighted) diagram each site carries an additive term, with a similar bending effect. The highest-exposure rule of this paper is not literally a multiplicatively weighted diagram, because exposure severity is a banded function of distance and charge rather than a single scalar weight, but the weighted-Voronoi picture is the correct intuition for it: a larger charge behaves like a more strongly weighted site and claims a region of positions that are physically closer to a smaller competitor. We return to this in Section 4. This material is standard computational-geometry background and is presented without attribution to a specific source.

2.3 Relationship to the exposure model

The selection rules in this paper are not new physics. They are decision rules layered on the exposure model of Paper 2, which in turn rests on the cordon model of Paper 1 and the underlying blast and fragment correlations (Baker et al., 1983; Brode, 1955; Department of Defense Explosives Safety Board, 2009; Kingery & Bulmash, 1984; Mills, 1987). The contribution here is to use that exposure model as the cost function of a spatial partition, so that the partition reflects danger rather than distance.

3 Model

3.1 Incident set and per-incident quantities

Let an incident set $\{I_k\}_{k=1}^{K}$ be given. Each I_k carries a location and the cordon and exposure parameters produced by Paper 1, namely the inner and outer cordon radii, the maximum and hazardous fragment distances, and the TNT-equivalent charge. For a responder at position x, define two per-incident quantities. The first is the great-circle distance

$$d_k(x) = \operatorname{dist}(x, I_k)\,, \tag{1}$$

computed by the haversine formula. The second is the exposure severity

$$e_k(x) = \operatorname{level}\big(\operatorname{exposure}(x, I_k)\big) \in \{0, 1, 2, 3, 4\}, \tag{2}$$

the integer band index returned by the exposure engine of Paper 2, where 0 is clear, 1 cordon, 2 caution, 3 danger, and 4 lethal. The severity folds together overpressure band and fragment-hit probability, so it already expresses the danger that a distance alone cannot.

3.2 Single-responder strategies

The system offers two strategies for selecting the active incident for one responder,

$$\text{nearest:} \quad k^* = \arg\min_k \; d_k(x), \tag{3}$$

$$\text{highest exposure:} \quad k^* = \arg\max_k \; e_k(x). \tag{4}$$

Each uses the other quantity as a tie-break. The nearest strategy breaks ties toward higher exposure, and the highest-exposure strategy breaks ties toward nearer distance, so that both are fully determined. The highest-exposure strategy is the default. A responder should be shown the incident most able to harm them where they stand, and only when two incidents are equally dangerous does proximity decide. The nearest strategy is retained because it is useful when all charges are comparable, when the responder explicitly wants the closest item, or as a transparent baseline against which the exposure-aware behaviour can be checked.

3.3 Team strategies

When several responders are present at known positions, two team strategies are available. Independent assignment applies the single-responder rule of choice to each responder separately, so two responders near the same dominant charge both receive it. This is correct when responders are converging on one problem and all need the same picture.

The distinct strategy instead divides the problem. It assigns at most one responder per incident by forming all responder-incident pairs, sorting them by ascending distance, and greedily accepting the shortest pair whose responder and incident are both still unassigned, until no further pair can be accepted. Its objective is to minimise the total of the accepted nearest-pairing distances. We note plainly that this greedy construction is not guaranteed to find the globally optimal matching under that objective, and that it does not optimise an exposure-weighted cost; Section 7 discusses the optimal formulation. Surplus responders, when there are more responders than incidents, are left unassigned rather than doubled onto an already-covered incident.

3.4 Serialised output

Whichever strategy is selected, the result is a chosen incident together with its index, the responder-incident distance, and the exposure severity at the responder's position. This result, together with the full incident list, is what the companion application serialises into a compact versioned record and pushes to the wrist-worn display, so that the watch renders the cordon for the incident the model selected.

4 Results

4.1 Two incidents

Figure 1 shows the assignment field for two incidents, a small 5 kg charge on the left and a large 800 kg charge on the right, across a grid of responder positions. Under the nearest rule the plane is split by the perpendicular bisector of the segment joining the two incidents, a straight vertical midline that takes no account of charge size. This is the Voronoi partition of the two points. Under the highest-exposure rule the boundary bows toward the small charge. A band of positions that are physically closer to the small charge are assigned to the large one, because the large charge governs their exposure there. This is the weighted-Voronoi behaviour anticipated in Section 2: the larger charge acts as a more strongly weighted site and claims territory across the geometric midline. The bowed region is precisely the set of positions where a nearest rule would direct a responder to watch the lesser hazard.

4.2 Worked example

Table 1 makes the disagreement concrete at three representative positions, using the two-incident scene of Figure 1 with the small charge at the left marker and the large charge at the right marker. At a position near the small charge but well within the large charge's hazard footprint, the two strategies disagree: nearest selects the small charge, highest exposure selects the large one. At positions clearly dominated by one incident, the two strategies agree. The disagreement is not an edge case to be smoothed away. It is the entire reason the exposure-aware rule exists.

4.3 Three incidents

The partition generalises beyond two competitors. Figure 2 shows the highest-exposure assignment for three simultaneous incidents of small, medium, and large charge. Each incident claims a contiguous region of responder positions, and the boundaries between regions are curved rather than straight, with the large charge claiming the most territory for a given separation. The

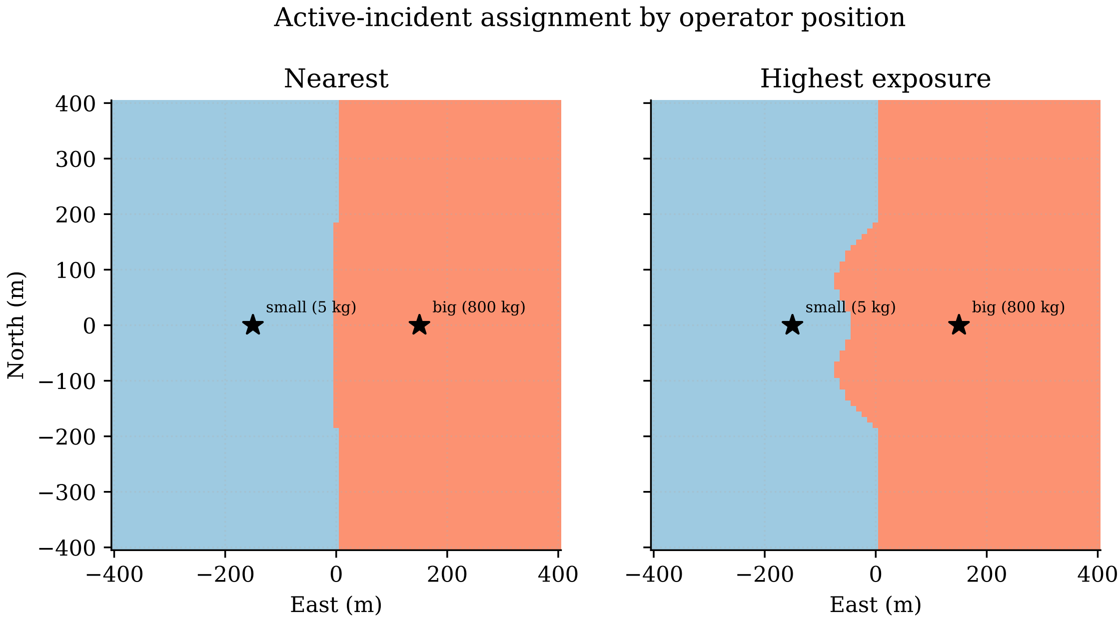

Figure 1. Active-incident assignment by responder position for a small (left) and a large (right) charge, generated from the shipped assignment and exposure code. Left panel: nearest, a straight Voronoi midline. Right panel: highest exposure, whose boundary bends toward the small charge and reassigns a band of nearby positions to the more dangerous large charge.

Table 1. Worked selection at three responder positions for the two-incident scene. Distances d_k and severities e_k are illustrative of the regimes; S denotes the small charge, L the large charge. The chosen incident differs between strategies only where a nearer small charge competes with a more dangerous large one.

	d_S (m)	d_L (m)	e_S	e_L	nearest	highest exposure
Near small, inside large footprint	30	170	1	2	S	L
Deep inside small only	8	320	3	1	S	S
Deep inside large only	260	25	0	4	L	L

result is a clean partition of the operating area into zones of attention, one per incident, that reflects danger rather than mere proximity.

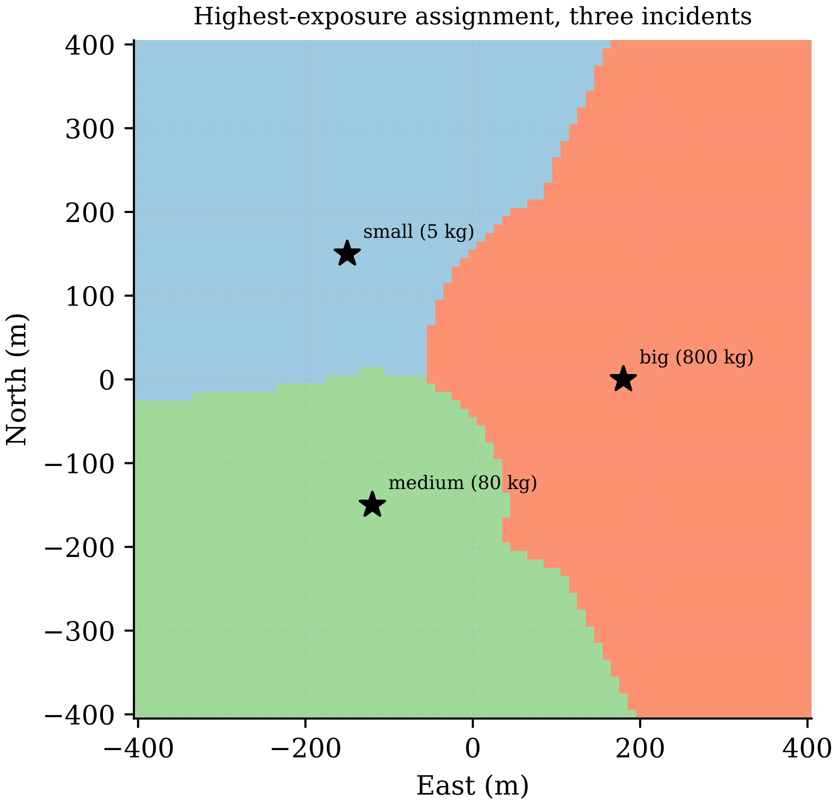

Figure 2. Highest-exposure assignment for three simultaneous incidents of small, medium, and large charge, generated from the shipped code. Each incident claims a contiguous region; boundaries curve in proportion to charge size, the weighted-Voronoi behaviour seen with two incidents.

5 Implementation and Traceability

The assignment functions live in the same shared library as the cordon and exposure models, and they call the exposure engine directly to obtain $e_k(x)$. The decision maps in Figures 1 and 2 were produced by evaluating the shipped assignment functions over a grid of responder positions and recording the selected incident at each cell. The decision boundary drawn in each figure is therefore the decision boundary the product computes, not a separate re-derivation. The selection that these functions return is the selection the companion application serialises and pushes to the wrist-worn display, so the chain from charge parameters to the cordon shown on the responder's wrist passes through exactly the code these figures exercise.

6 Limitations and Threats to Validity

Exposure severity is a banded integer rather than a continuous field, so the highest-exposure boundary is the contour of a step function and is slightly ragged at grid resolution. This is visible as the stepped edges in the decision maps. A continuous severity score would yield smoother boundaries, at the cost of introducing a tuning choice that the banded scheme avoids.

The distinct team strategy is greedy. It minimises the total of accepted nearest-pairing distances by a shortest-pair-first heuristic, which is fast and predictable but is not guaranteed to attain the global optimum of that objective, and it does not optimise an exposure-weighted cost at all.

Finally, and most importantly, every strategy decides which incident a display presents. None of them decides whether a responder should move, where they should stand, or how they should proceed. The model informs attention; it does not direct action, and it does not replace command explosive-safety authority.

7 Future Work

Two extensions follow naturally. The distinct team strategy could be posed as a formal assignment problem with an exposure-weighted cost and solved to optimality, trading a modest increase in computation for a guaranteed best matching. Second, the present treatment is static: it assigns at a position. A dynamic formulation would re-assign as responders move and as incidents are added or resolved, with hysteresis to avoid rapid switching of the active incident near a partition boundary. Both extensions preserve the central design principle of this series, that selection and readout are computed from one shared exposure model.

8 Conclusion

When hazards are simultaneous, the right cordon to show is the one for the hazard most able to harm the responder, which is not always the nearest. By defining exposure-aware selection rules on top of the shared cordon and exposure engine, by interpreting their behaviour through weighted Voronoi geometry, and by showing that the highest-exposure boundary departs from the nearest midline exactly where a larger charge dominates, this paper completes a single, traceable chain. Charge parameters become a cordon in Paper 1, the cordon becomes live exposure in Paper 2, and live exposure becomes the assignment of attention in this paper, all the way to the incident shown on a responder's wrist. The three papers together document the calculations behind The Long Walk system as one auditable whole.

References

Baker, W. E., Cox, P. A., Westine, P. S., Kulesz, J. J., & Strehlow, R. A. (1983). *Explosion hazards and evaluation*. Elsevier.

Bowen, I. G., Fletcher, E. R., & Richmond, D. R. (1968). *Estimate of man's tolerance to the direct effects of air blast* (tech. rep. No. DASA-2113). Defense Atomic Support Agency.

Brode, H. L. (1955). Numerical solutions of spherical blast waves. *Journal of Applied Physics*, *26*(6), 766–775. https://doi.org/10.1063/1.1722085

Department of Defense Explosives Safety Board. (2009). *Fragment and debris hazards* (tech. rep. No. Technical Paper No. 12). Department of Defense Explosives Safety Board.

Kingery, C. N., & Bulmash, G. (1984). *Airblast parameters from TNT spherical air burst and hemispherical surface burst* (tech. rep. No. ARBRL-TR-02555). U.S. Army Ballistic Research Laboratory. Aberdeen Proving Ground, MD.

Mills, C. A. (1987). The design of concrete structures to resist explosions and weapon effects. *Proceedings of the 1st International Conference on Concrete for Hazard Protection*.

U.S. Department of Defense. (2008). *Structures to resist the effects of accidental explosions* (tech. rep. No. UFC 3-340-02). Unified Facilities Criteria.

U.S. Department of Defense. (2019). *DoD ammunition and explosives safety standards* (tech. rep. No. DESR 6055.09). Defense Explosives Safety Regulation.

U.S. Department of Homeland Security & Bureau of Alcohol, Tobacco, Firearms and Explosives. (2016). *DHS/ATF bomb threat standoff card* (tech. rep.). U.S. Department of Homeland Security.