THE COGNITIVE SPACE OF AIR TRAFFIC CONTROL: A PARAMETRIC DYNAMIC TOPOLOGICAL MODEL[1]

Steven Cushing

Boston University
Computer Science B-4
755 Commonwealth Avenue
Boston, MA 02215

Abstract

Recent observational work of controller behavior in simulations of air traffic control sessions suggests that controllers formulate and modify their plans in terms of *clusters* of aircraft, rather than *individual* aircraft, and that they cluster aircraft based on their *closeness in an abstract cognitive space*, rather than simple *separation in physical space*. A mathematical model of that space is presented as a background for further work to determine the *cognitive strategies* that controllers use to navigate that space. The model is *topological* in that *neighborhood constraints* play a central role; it is *dynamic* in that *more than one topology interact* to define its essential characteristics; and it is *parametric* in that *an entire class of spaces can be obtained* by varying the values of some parameters.

With the presented model as background, some hypotheses as to controller strategies are suggested and examples given to illustrate them. For example, controllers appear to segment their work into *episodes* defined in terms of the interactions of clusters and to *prioritize* the sub-tasks within these episodes, with different strategies for different sub-tasks. Some specific questions suggested by the hypotheses are raised, and some theoretical and practical implications are pointed out. For example, controllers appear to change their plans *consequent upon* changes in perceived clustering: *deliberate* cognitive acts are triggered by *presented* changes in conceptualization. This has implications for tool development, in that it underscores the need to limit the extent to which automated aids should be allowed to deviate from actual controller practice.

---

[1] Supported, in part, by Contract No. NGT 47-003-029, NASA-Langley Research Center, Hampton, VA. I would like to thank Herb Armstrong, Hugh Bergeron, Greg Bonadies, Randy Harris, Gary Lohr, Renate Rofske-Hofstrand, and George Steinmetz for invaluable assistance on the empirical work that underlies the theoretical model presented here. I also thank Eric Braude, Robert Kuhns, and especially Carol Munroe for very helpful comments on an earlier draft.

1. Background. Recent observational work on controller behavior in air traffic control simulation sessions suggests that controllers formulate and modify their plans in terms of *clusters* of aircraft, rather than *individual* aircraft, and that they cluster aircraft based on their *closeness in an abstract cognitive space*, rather than simple *separation in physical space* (Cushing, 1989; also see Bregman et al., 1988; Cushing, 1990; NASA/OAST, 1988; Roske-Hofstrand, 1988; Roske-Hofstrand et al., 1989; and Wesson, 1977; for further background and discussion). In this paper a mathematical model of that space is presented based on that work and as background for further work to determine the *cognitive strategies* that controllers use to navigate that space. The model is *topological* in that *neighborhood constraints* play a central role; it is *dynamic* in that *more than one topology interact* to define its essential characteristics; and it is *parametric* in that *an entire class of spaces can be obtained* by varying the values of some *external parameters.* The several topologies are also distinguished by the values of *internal parameters*. The parameters are presented in **Section 2** and the model itself in **Section 3**. With it as background, some hypotheses as to controller strategies are suggested in **Section 4**, along with some examples that illustrate them and some specific questions for further investigation.

2. The Parameters. The essence of the model is illustrated in Figures 1-9. Arriving aircraft are identified with (an initial segment of) the set $\mathbf{Z}^+$ of positive integers, and the controller's task is modelled as the construction of a permutation λ on $\mathbf{Z}^+$. In other words, aircraft that are presented in an order of arrival must be safely rearranged into an order for landing. The construction of an appropriate λ is subject to constraints that are central to the structure of the space. Aircraft that are within a neighborhood of each other in *arrival order* must be kept within a neighborhood of each other in *landing order*, with the respective neighborhoods defined by the parameters $\pi_{\mathbf{a}}$ and $\pi_{\mathbf{l}}$. Similarly, aircraft that are within a neighborhood of each other in *arrival time* must be kept within a neighborhood of each other in *landing time*, with the respective neighborhoods defined by the parameters $\tau_{\mathbf{a}}$ and $\tau_{\mathbf{l}}$. The topologies that are determined by these neighborhood constraints partially characterize the structure of the space.

Neighborhood constraints defined in terms of six other parameters determine further topologies that complete that characterization. Four of these, like $\pi_{\mathbf{a}}$ and $\pi_{\mathbf{l}}$, are *numerical parameters*: the ***minimum separation*** (i.e., physical distance) σ that must be maintained at all times between any two aircraft, the ***maximum order difference*** ω and ***maximum cognitive distance*** κ for aircraft to be in a cluster, and the ***minimum duration interval*** ι that an aircraft can be in a cluster. In actual practice, σ has the value 3.0 (miles) for good weather and 5.0 for bad. A possible refinement would be to allow a further parameter $\sigma_{\mathbf{h}}$ for "heavy" aircraft, which require greater separation, but this can be avoided if the two separation values are taken to be systematically related, e.g., by a constant difference or ratio. The value of ι depends on short-term memory capacity and perception thresholds and can thus be expected, like ω and κ, to vary from controller to controller, in contrast to the values of σ, which are set by regulation.

Finally, we have two *functional parameters*: the ***physical distance*** ϕ and the ***cognitive distance*** χ . We can reasonably assume that ϕ is unique and fully understood, but determining the character of χ is intricately intertwined with the related problem of determining the cognitive strategies in which it is used and thus, other than being required to be a *distance metric*, is beyond the scope of the present paper.[2] Any of the parameters can be varied at will for purely investigative purposes, i.e., modeling, simulation, or the like.

3. The Model. Since we identify aircraft with their positions in arrival and landing order, i.e., before and after permutation by the controller, the parameters that bound their neighborhoods and clusters are naturally taken also to be integers. Physical or cognitive distances or times, however, are more reasonably taken to be real-valued. We thus begin as follows:

Stipulation: Numerical Parameters. Choose fixed values

$\pi_a, \pi_l, \omega \in \mathbf{Z^+}$,

$\tau_a, \tau_l, \sigma, \iota, \kappa \in \mathbf{R^+}$.

Safety requires, first and foremost, that the *physical distance* between aircraft not be permitted to attain a value less than a specified minimum, identified above as the ***minimum separation*** σ . It is apparent from controller behavior, however, that a very different measure of *cognitive distance* also plays a role in determining which aircraft are taken as being "close enough" to be considered together as a ***cluster***, as illustrated in Figures 4 and 5. Even aircraft whose arrival occurs at opposite ends of the controller's screen and are thus widely separated physically can be grouped together in planning, as the required permutation is constructed. We thus continue as follows:

Stipulation: Functional Parameters. Choose distinct fixed values

$\phi, \chi : \mathbf{Z^+} \times \mathbf{Z^+} \times \mathbf{R^+} \rightarrow \mathbf{R}$,

subject to the following constraint:

For $\delta \in \{\phi, \chi\}$, $\forall\ p_1, p_2, p_3 \in \mathbf{Z^+}$, $t \in \mathbf{R^+}$,

$0 \leq \delta(p_1, p_2, t)$,

$\delta(p_1, p_2, t) = \delta(p_2, p_1, t)$,

$\delta(p_1, p_2, t) \leq \delta(p_1, p_3, t) + \delta(p_3, p_2, t)$.

[2] Compare the analogous difference, in physics, between claiming that the universe is a four–dimensional differentiable manifold and claiming that it is a *particular* such manifold or, in linguistics, between claiming that language is characterized by an underlying autonomous universal syntax and claiming that that syntax is embodied in grammatical rules or principles *of a particular form*. Though equally non-trivial, the two sorts of claims differ substantially in strength, the first providing the background for investigating the second, further refinements of which can eventually lead to the confirmation of both (or not). The space proposed here embodies a claim of the former sort; further characterizing χ would involve successive refinements of claims of the latter.

This is the standard definition of ***distance metric*** in a metric space, but *relativized* to times and with an *internal parameter* δ that takes the *two* functions ϕ and χ as instances, thereby defining two different topologies on the space. Note that the stipulated relativization allows δ itself to vary with time, with different results possible at different times for the same ordering or spatial arrangement of aircraft. Though not sufficiently constraining for *physical* distance, this flexibility is desirable for *cognitive* distance, which can be expected to vary with time of day, work load, stress, and the like, as these affect the controller's mental state.

The *difference in arrival order,* $p_2 - p_1$, is a further distance metric, *not* dependent on time, that, along with χ, constrains the formation of clusters. *It is the occurrence of clustering that makes controller cognition a non-trivial problem.* If aircraft were treated entirely as individuals, it would be a simple matter, at least in principle, to automate the process, by assigning each aircraft a dedicated processor to maintain separation σ from every other, as is apparent from the Figures. With each aircraft capable of fending for itself, air traffic control would be rendered superfluous. Clustering reflects the need, however, *to plan ahead* in maintaining σ, not just *now*, but *throughout* the permutation process. It is thus in the internal logic of clusters that we can most reasonably expect to find indications of how this planning proceeds.

Definition: **Neighborhood.** Let $\mathbf{D} \in \{\mathbf{Z^+}, \mathbf{R^+}\}$. $\forall\ d_1, d_2, \beta \in \mathbf{D}$, d_2 **is in a** β**–neighborhood of** d_1, $\mathbf{N}_\beta(d_1, d_2)$, if $|d_2 - d_1| < \beta$.

Corollary: (1) **Reflexive.** $\forall\ d, \beta \in \mathbf{D}, \mathbf{N}_\beta(d, d)$.

(2) **Symmetric.** $\forall\ d_1, d_2, \beta \in \mathbf{D}, \mathbf{N}_\beta(d_1, d_2) \Rightarrow \mathbf{N}_\beta(d_2, d_1)$.

(3) **Intransitive.** $\exists\ d_1, d_2, d_3, \beta \in \mathbf{D}, \mathbf{N}_\beta(d_1, d_2) \wedge \mathbf{N}_\beta(d_2, d_3) \wedge \neg\mathbf{N}_\beta(d_1, d_3)$.

Again, this is a standard definition of *neighborhood* in a metric space, but with the ***domain*** **D** and the ***neighborhood bound*** β as parameters *internal to the model*, taking $\mathbf{Z^+}$ or $\mathbf{R^+}$ and the external parameters π_a, π_l, τ_a, τ_l, or ω as values, respectively, in various places in the model. The corollary is also standard and serves mainly to disqualify $\mathbf{N}_\beta$ from being an *equivalence relation.*

Definition: Clustering. A function $\mathbf{c}: \mathbf{Z^+} \times \mathbf{R^+} \rightarrow 2^{\mathbf{Z^+}}$ is a $(\omega, \iota, \kappa, \chi)$**–clustering** if it satisfies the following constraints:

(1) (a) $\forall\ p \in \mathbf{Z^+}, \exists\ t_{p,a}, t_{p,l} \in \mathbf{R^+}, \forall\ t \in \mathbf{R^+}$,

$[t \in [t_{p,a}, t_{p,l}] \Rightarrow p \in \mathbf{c}(p, t)\ ] \wedge [t \notin [t_{p,a}, t_{p,l}] \Rightarrow \mathbf{c}(p, t) = \varnothing\ ]$,

(b) $\forall\ p_1, p_2 \in \mathbf{Z^+}, \forall\ t \in \mathbf{R^+}, [p_2 \in \mathbf{c}(p_1, t) \Rightarrow p_1 \in \mathbf{c}(p_2, t)]$,

(c) $\forall\ p_1, p_2, p_3 \in \mathbf{Z^+}, \forall\ t \in \mathbf{R^+}$,

$[p_2 \in \mathbf{c}(p_1, t) \wedge p_3 \in \mathbf{c}(p_2, t)] \Rightarrow p_3 \in \mathbf{c}(p_1, t)$,

(2) (a) $\forall\ p_1, p_2 \in \mathbf{Z^+}, t \in \mathbf{R^+}, [p_2 \in \mathbf{c}(p_1, t) \Rightarrow \mathbf{N}_\omega(p_1, p_2)]$,

(b) $\forall\ p_1, p_2 \in \mathbf{Z^+}, t \in \mathbf{R^+}, [p_2 \in \mathbf{c}(p_1, t) \Rightarrow \chi(p_1, p_2, t) < \kappa]$,

(3) $\forall\ p \in \mathbf{Z^+}, t \in \mathbf{R^+}$,

$$[\mathbf{c}(p, t) \neq \varnothing \Rightarrow [\exists\ t_1, t_2 \in \mathbf{R^+}, t_2 - t_1 \geq \iota,\ t \in [t_1, t_2],$$
$$\forall\ t' \in (t_1, t_2), [\mathbf{c}(p, t') \neq \varnothing \Rightarrow \mathbf{c}(p, t') = \mathbf{c}(p, t)]]].$$

A *clustering* assigns to *each* aircraft (in $\mathbf{Z^+}$) a *set* of aircraft, its *cluster*, (in $2^{\mathbf{Z^+}}$) that can *change* as time (in $\mathbf{R^+}$) passes, subject to three constraints. Constraint **(1)** says that ***being–clustered–with***, unlike ***being–in–a–neighborhood–of***, is an *equivalence relation* at a time t on sets of aircraft with overlapping ***salience intervals*** $[t_{p,a}, t_{p,l}]$: an aircraft is always in its own cluster between its arrival time and its landing time, but has *no* genuine cluster (i.e., its cluster = $\varnothing$) otherwise. Outside that interval, an aircraft is of no cognitive interest, in the present context; it need not be worried about before it arrives and after it lands. Assigning it a cluster of $\varnothing$ simplifies the formulations by avoiding the need for partial functions. An aircraft can be viewed as *comprising* its own cluster whenever it is treated individually. Constraint **(2)** relates clustering to the two *non-physical distance metrics* of the model by imposing *proximity constraints*: one aircraft can be in the cluster of another aircraft *only* if they are within ω of each other in ***arrival order*** and within κ of each other in ***cognitive distance***. The actual values of ω and κ will depend on the individual controller, but they can be expected to be intimately connected to limits on short–term memory. *Characterizing* cognitive distance amounts to *solving simultaneously the conditions of the definition* for suitable functions χ .[3] Constraint **(3)** says that a clustering can assign a particular set of aircraft as the cluster of a particular aircraft at a salient time *only* if it assigns that set as the cluster of that aircraft *within an interval that contains that time*, of width no less than ι. This guarantees *a degree of* (i.e., step–wise) *continuity* so that *changes* in clustering remain realistic, without diminishing the flexibility that makes those changes useful.

The parameters ω and ι serve to define upper bounds on the size and the number of clusters and on the number of times they can change, as follows:

Theorem: Let C be the set of all $(\omega, \iota, \kappa, \chi)$–clusterings.

(1) $\forall\ p \in \mathbf{Z^+}, t \in \mathbf{R^+}, \mathbf{c} \in C$,

$|\mathbf{c}(p, t)| \leq 2\omega - 1$.

(2) $\forall\ p \in \mathbf{Z^+}, t \in [t_{p,a}, t_{p,l}]$,

$|\{\mathbf{c}(p, t) \mid \mathbf{c} \in C\}| \leq 2^{2(\omega - 1)}$.

(3) $\forall\ p \in \mathbf{Z^+}, \mathbf{c} \in C$,

$$|\{t \mid t \in [t_{p,a}, t_{p,l}),\ \exists\ \varepsilon \in \mathbf{R^+}, \mathbf{c}(p, t - \varepsilon) \neq \mathbf{c}(p, t + \varepsilon)\}|$$
$$< ((t_{p,l} - t_{p,a}) / \iota) + 2.$$

---

[3] See note 2.

Proof: (1) Planes available for p's cluster range from $p - (\omega - 1)$ through $p + (\omega - 1)$, including p itself:

$\therefore\ (\omega - 1) + (\omega - 1) + 1 = 2\omega \quad 1.$

(2) Clusters generally available for p include *all subsets* of the set of planes in (1) *that contain p*:

$\therefore\ (1/2)(2^{2\omega \quad 1}) = 2^{2\omega \ - 2} = 2^{2(\omega - 1)}.$

The ≤ is required, rather than =, because of the current indeterminacy of χ, which can be expected, when determined, to rule out possible clusters that are permitted by the other constraints, and also, independently of χ, because of the fact that two different planes will not have *identical* salience intervals (see ***landing sequence*** below).

(3) A plane's cluster becomes non-null upon arrival and can then change at any time thereafter. After the first change, each cluster must then persist for an interval of at least ι in duration. Let $L = (t_{p,l} - t_{p,a})$ and assume each interval has duration exactly ι. If ι exactly divides L, there are L / ι change points if $t_{p,a}$ begins an interval of length ι, and $(L / \iota) + 1$ otherwise. If ι does not exactly divide L, there are $[(L / \iota) + 1]$ change points if $t_{p,a}$ begins an interval of length ι, and $[(L / \iota) + 2]$ otherwise, where $[x]$ is the greatest integer in x. If one or more intervals has duration greater than ι, the number of change points does not increase.

In other words, the number *of distinct aircraft* that are *in the cluster of a particular aircraft at a particular time* is bounded above by $2\omega - 1$; the number *of distinct potential clusters* that are available to be associated *with a particular aircraft at a particular time, between the time it arrives and the time it lands,* is bounded above by $2^{2(\omega - 1)}$; and the number *of distinct times* that a controller *assigns a new* ***non-null*** (hence the ")" in $[t_{p,a}, t_{p,l})$ here) *cluster to a particular aircraft* is bounded above by $((t_{p,l} - t_{p,a}) / \iota) + 2$.

Finally, we can characterize ***the controller's task*** as follows:

Definition: Landing Sequence. A permutation λ on $\mathbf{Z^+}$, i.e., a bijection $\lambda : \mathbf{Z^+} \rightarrow \mathbf{Z^+}$, is a $(\pi_a, \pi_l, \tau_a, \tau_l)$**–landing sequence** if it satisfies the following constraint:

$$\forall\ p_1, p_2 \in \mathbf{Z^+},$$
$$[\,[N\pi_a(p_1, p_2) \Rightarrow N\pi_l(\lambda(p_1), \lambda(p_2))]$$
$$\wedge\ [N\tau_a(t_{p_1,a}, t_{p_2,a}) \Rightarrow N\tau_l(t_{p_1,l}, t_{p_2,l})]$$
$$\wedge\ [p_1 \neq p_2 \Rightarrow t_{p_1,l} \neq t_{p_2,l}]\,].$$

Controller's Task: Construct a clustering and a landing sequence subject to the following constraint:

$$\forall\ p_1, p_2 \in \mathbf{Z^+},\ t \in [t_{p_1,a}, t_{p_1,l}] \cap [t_{p_2,a}, t_{p_2,l}],\quad \phi(p_1, p_2, t) \geq \sigma\ .$$

In other words, aircraft that are "near" each other in *arrival* order or time must be "near" each other in *landing* order or time, respectively, though the "nearness" *criteria* for arrival and landing need not be the same and the landing *times* must be different. Furthermore, no matter how they get permuted or what clustering is used, two aircraft must never be permitted to approach each other to within a physical distance of σ. Constructing a landing sequence is the *explicit* requirement of the controller's task; the need to construct a clustering arises from the inherent complexity of that task in the context of human cognition.

4. Further Work In Progress. The above model provides the *general cognitive framework* within which controllers appear to do their work. The next task is to determine the cognitive *strategies* that controllers use to navigate the space the model characterizes. Work thus far has suggested the following hypotheses (Cushing, 1989):

(1) Controllers segment their work temporally and dynamically into (*sometimes overlapping*) *episodes* and *sub-episodes* defined in terms of the interactions of aircraft clusters;
(2) Controllers *prioritize* the sub-tasks within their episodes, with different strategies for different sub-tasks; and
(3) Controllers change plans *consequent upon* changes in perceived clustering: *deliberate* cognitive acts are triggered by *presented* changes in conceptualization.

Hypothesis (1) is illustrated throughout the Figures. As an example of hypothesis (2), a controller in a simulation session checks that separation of aircraft is adequate both before and after doing an artificial "side–task" consisting of reading extraneous information about the weather, but is less thorough in checking before scanning to accept responsibility for an aircraft that is being handed off to him by another controller. This suggests that he considers the latter task to be more important and in need of more immediate attention when it arises. Hypothesis (3) has particular implications for tool development, in that it underscores the need to limit the extent to which *automated aids*, such as expert system or decision support tools, should be allowed to deviate from actual controller practice. Further work is needed to quantify and test these hypotheses.

In particular, the following questions need to be answered: (1) What factors *other than arrival order* play a role in clustering?[4] (2) Does the controller check aircraft separation *in preparation* for doing the "side-task," or does he do the "side-task" *after having checked* separation? (3) To what extent does the controller maintain separation *of clusters*, and to what extent is he willing to *shuffle* (i.e., modify and mix) them? (4) How often and why does the controller scan *back* to aircraft that are *already* lined up for landing, while focusing primarily on a *later* cluster, and how often and why does he scan to *outliers* beginning a *new* cluster, while focusing primarily on an earlier one?

It is anticipated that the pursuit of answers to these and related questions will lead to refinements in the above model, as the strategies it supports are unraveled. It will be necessary to investigate the extent to which controllers *differ* in their choice of strategies, both from each other and from non–controller control subjects, in order to determine the extent to which the strategies used are *learned*, rather than corollaries of *inherent* properties of human perceptual and cognitive mechanisms. How this question gets answered has implications for *training methods*, even aside from the development of support tools for assisting controllers on the job.

---

[4] I.e., just what is it that distinguishes **(b)** from **(a)** in condition **(2)** of the clustering definition?

References

Bregman, Howard L.; McCabe, Warren L.; and Sutcliffe, William G. 1988. "Capturing Air Traffic Controller Expertise for Incorporation in Automated Air Traffic Control Systems." ***Proceedings of the Human Factors Society, 32nd Annual Meeting.*** pp. 1031-1035.

Cushing, Steven. 1989. "From Where They Look to What They Think: Determining Controllers' Cognitive Strategies from Oculometer Scanning Data." In Tiwari, Surendra N. (Comp.). NASA Contractor Report 181894. Contract No. NGT 47-003-029. NASA-Langley Research Center, Hampton, VA. September 1989.

Cushing, Steven. 1990. Letter. ***The Sciences.*** 30(1):14. January/February 1990.

NASA/OAST. 1988. "Aviation Safety/Automation: A Plan for the Research Initiative." National Aeronautics and Space Administration, Office of Aeronautics and Space Technology. November, 1988.

Roske-Hofstrand, Renate. 1988. "Discourse Understanding of Displayed Actions: An Investigation of the Temporal Relations of Event Schemas and their Effect on Reconstruction." Doctoral Dissertation. New Mexico State University, Las Cruces.

Roske-Hofstrand, Renate; Armstrong, Herbert; and Bergeron, Hugh. 1989. "Applied Cognitive Research in Air Traffic Control: A Concept Paper and Research Plan." Draft Report. NASA-Langley Research Center, Hampton, VA.

Wesson, Robert Bell. 1977. "Problem-solving with Simulation in the World of an Air Traffic Controller." Doctoral Dissertation. University of Texas, Austin. N78–27098. University Microfilms Dissertation Information Service.

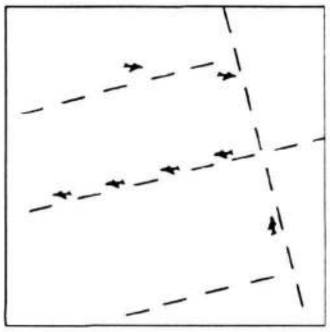

Figure 1: A typical ATC *arrival* situation: Aircraft arriving from *several directions* must be *lined up* for landing.

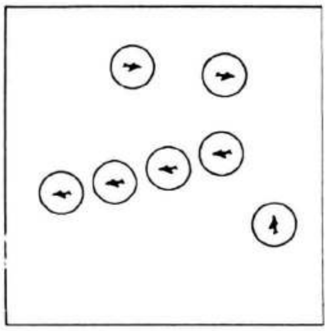

Figure 2: *Minimum separation* must be maintained between *any two aircraft* at *any time.*

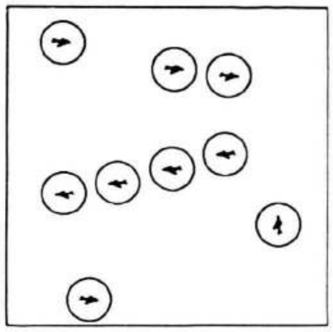

Figure 3: *Distant, but contemporaneous,* arrivals are treated together in *planning.*

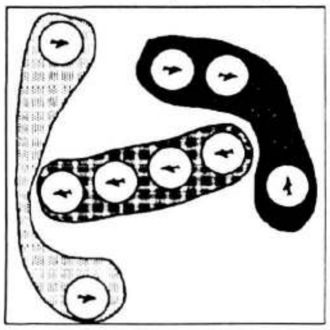

Figure 4: Aircraft are *clustered* by "closeness" in an abstract *cognitive* space, rather than in actual *physical* space.

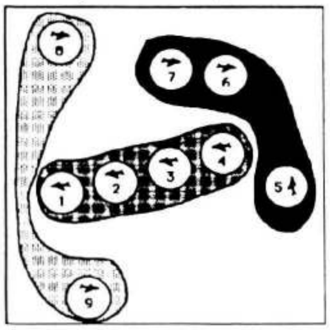

Figure 5: *Cognitive "closeness"* is measured by a *spatio-temporal metric* distinct from *arrival order* and *physical separation.*

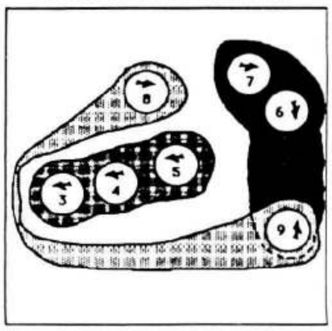

Figure 6: Clusterings can *change,* due to *unforeseen circumstances*: Aircraft can *leave* one cluster and *enter* another.

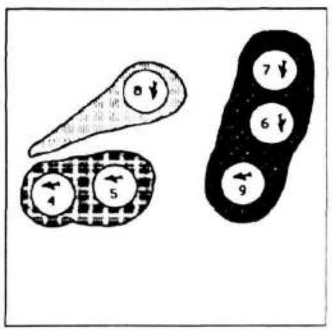

Figure 7: ATC space has a *dynamic topology* that reflects the *interaction* of *multiple factors.*

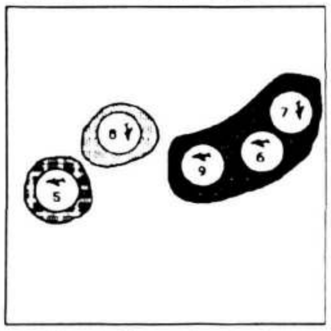

Figure 8: Changes in *clustering* correlate with changes in *plan,* in *response* to new circumstances.

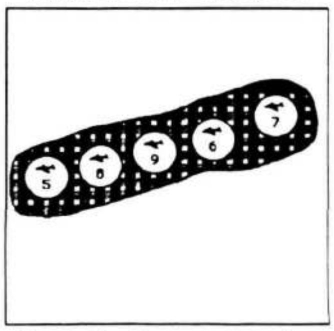

Figure 9: Aircraft *merge* on final to form a single cluster: The *landing* order is a *permutation* of the *arrival* order.