Viewpoint

How Well Do We Understand Air Command and Control?

By Air Commodore Alistair Byford

Introduction

Traditionally, the UK has exercised air command and control (C2) through a model of centralised control, decentralised execution, enabled by the concept of mission command. However, our most recent statement of environmental doctrine, *JDP 0-30 'UK Air and Space Doctrine'*,[1] mandates a change to a more flexible approach, *centralised control, adaptive execution*. In my experience, airmen and airwomen are not particularly good at understanding C2 or giving it the prominence it deserves. And this is problematic, because C2 is absolutely fundamental to the effective employment of air power. So I think it is worth considering how air C2 has evolved, why the time is ripe for change and what the implications are likely to be.

Command versus Control

We need to begin by establishing some principles and defining the terms. *British Defence Doctrine*[2] tells us that we vest *command* of a force element, organisation or operation in an individual to assign direct authority, responsibility and accountability. This differs markedly from *control*, which is merely 'the coordination of the activities necessary to achieve the commander's intent'. Put simply, command is exercised by a commander, while control is a staff-led function. This means command is inherently a human activity in a way that control is not, because we can hold a commander - as an individual person - to account for his or her actions, while control may be more about processes and procedures.

Both command and control can be facilitated by common doctrine and communication systems that allow us to share information quickly and easily. However, the peculiar nature of

air power and the specific mechanisms we use to deliver it - think, for example, of the Airspace Coordination Order (ACO) and the Air Tasking Order (ATO) - mean that in my opinion, airmen and airwomen are particularly prone to being seduced by the technicalities of control at the expense of a proper emphasis on the key principles of command. This matters, because it is the air commander's personal attributes and qualities of command (in terms of leadership, initiative, imagination, knowledge, judgement and professionalism) that will ultimately determine the success or failure of an air operation, not the sophistication of the systems and processes of control they use to achieve their intent; so ultimately, control is (or at least should be) just an enabler.

Centralised Control

This is not to say that control is unimportant. Centralised control emerged from the earliest days of military aviation as the key to the effective application of air power and this still remains the case today. It also explains why the Royal Air Force was created as the world's first independent air force nearly a century ago. In 1917, the Smuts Report recognised that decentralising control of air assets to the organic air components of the Army and Navy was inefficient and unsustainable, because it had dissipated the UK's overall air effort. Understandably, the Royal Flying Corps and Royal Naval Air Service had concentrated on direct support to the army in the field and the fleet at sea respectively. This made it impossible for them to allot sufficient resources to different and higher priorities, such as the air defence of the UK, and this became apparent in the wholly inadequate response to Zeppelin and long-range bomber raids. The solution was to centralise control of air assets and resources within a single command structure focused solely on the delivery of air power, leading to the creation of the Royal Air Force on 1 April 1918.

So centralising control is vital, because it assures unity of both purpose and effort. It means we can guarantee that inevitably scarce air resources are allocated and apportioned as efficiently and effectively as possible, meeting the commander's intent in line with the priorities he or she has set. Centralised control also embodies the principle of *strategy to task*, ensuring every sortie we fly can be tracked back to the strategic objectives set by higher authority. In contrast, decentralising control is invariably disastrous. In France in 1940, *L'Armée de l'Air* fighter squadrons were allocated to (and controlled by) individual army formations. This meant that despite overall numerical parity, 'penny-packets' of French fighters were spread along the front and simply could not match the concentrated force that the *Luftwaffe* was able to generate at the *schwerpunkts*, or points of main effort; the result was that the French air force was defeated in detail. Bizarrely, some French squadrons were held on the ground while crucial air battles were being fought nearby - sometimes in sight - because they were based on the wrong side of a divisional boundary or located in a different army formation's area of responsibility.

Execution: to centralise or not to centralise?

While centralised control is therefore a key principle - and there is no evidence to suggest this will change in the future, as air resources will, if anything, become even scarcer - the best method of executing tasks and missions has always been more open to question. At this point,

we should note that *execution* is shorthand for *execution authority* – the empowerment to take decisions, such as whether or not to release weapons, to abort or continue a sortie, or change a target – not the physical actions that can only be conducted in the cockpit or at the console, such as pulling a trigger or pressing a weapons release button. In simple terms, there are two options for execution: we can centralise it, so a remote headquarters or higher authority controls decision-making; or we can decentralise, sometimes right down to the operators of individual air assets.

Decentralised execution initially evolved not as a conscious, doctrinal preference, but rather as a pragmatic response to our limited ability to communicate with (and therefore control) aircraft in the air. During the First World War, ground-air communications were in their infancy, so centralised control was never a realistic option. But even towards the end of the Second World War, reliable, long-range, communication was difficult. So for example, Bomber Command Headquarters at RAF High Wycombe possessed neither the technological means nor the data-processing capacity to centralise the execution of raids comprising hundreds of bombers (all operating as separate speaking units) flying over Germany well beyond easy line-of-sight radio contact. Consequently, decentralising execution was the only practical method of exercising air C2 and squadrons were effectively tasked on the basis of "off you go, bomb Berlin, find your own way there and identify the target, drop your eggs, good luck, see you when you get back". However, it was clear that more control was necessary and Bomber Command soon introduced C2 measures to try to work round the constraints on communications. An Op Order was issued for every raid, detailing the targets, aiming points, bomb load, target marking technique and overall raid timings for each group, supported in turn by Group Op Orders specifying TOTs down to the minute; effectively an ATO as we would recognise it today. Also, the 'Master Bomber' concept was introduced later on in the war, where a mission commander would direct a raid in the air via R/T (radio-telephony). This was first practised by Wing Commander Guy Gibson over the Mohne Dam in 1943.

The evolution of C2 for Bomber Command indicates there will always be a desire (or requirement) to centralise aspects of execution and that the primacy of the decentralised approach to air C2 was never as cut and dry in practice as it now appears in retrospect. This is reinforced by the example of Fighter Command. The integrated air defence system used so effectively during the Battle of Britain in 1940 is often cited as a classic illustration of centralised control, decentralised execution in action. But arguably, this actually demonstrates that centralised execution (if technically viable) is actually a more effective way of exercising C2 in particular operational circumstances. It was the group and sector controllers, using the air picture provided by radio-direction finding stations (radar) and observer reports, who exercised decision-making authority from their centralised headquarters. They chose which squadrons to allocate to which raids; when to scramble them; where they should be positioned, in terms of height and geography; and if (and when) they should engage. The pilots were only responsible for the last stages of tactical execution after they had obtained visual contact with the enemy and been cleared to commit.

So what does this demonstrate about air C2? Simply that for Fighter Command in 1940, centralised execution was both possible and desirable. It was possible, because of the small number of speaking units involved[3] and easy, two-way, line-of-sight radio contact with fighters operating close to their sector control centres. And it was desirable, because the controllers were better informed than the aircrew and could therefore make better decisions; without the capability to datalink the air picture into the cockpit, only the controller on the ground could see if a raid was a feint or the main effort, assess whether an enemy formation was escorted or composed solely of fighters, and determine likely targets and the priority for their defence.

The Battle of Britain and Bomber Command's offensive against Germany therefore highlight the two key factors that shape the air C2 environment: first, our ability to communicate (including data transfer) will determine if centralised execution is feasible or not in the first place; and second, air C2 is most effective when decision-making authority is situated at the point where the best information is available. For Fighter Command, this was with the controller in the sector or group headquarters; but for Bomber Command, it had to be with the aircraft captain in the cockpit, because neither group nor command headquarters enjoyed access to any useful, real-time information once a raid was underway and the aircraft had crossed the coast.

Mission Command

Decentralised execution is closely tied to the notion of *mission command*. This concept, identified in *British Defence Doctrine* as a fundamental attribute of the 'British way of war',[4] is often misused or misunderstood. In fact, it was originally a German idea, and mission in this sense is simply a translation of *auftrag*, meaning a task directed towards fulfilling a particular purpose. Mission command deals with those aspects of leadership involved with setting and giving direction. It encapsulates the freedom for a subordinate to act (or not to act) on his or her own initiative, using allotted resources to achieve the higher commander's intent. This should be specified only in terms of what is to be achieved rather than the ways in which it is to be accomplished.

Mission command is attractive in an air power context, because it allows us to delegate the execution of certain command responsibilities, such as engagement or weapons-release authority, enabling tactical self-synchronisation and increasing tempo by shortening the decision cycle. This is particularly critical in the air environment, and it is no surprise that it was an airman, Colonel John Boyd, who articulated the concept of the Observe-Orientate-Decision-Action or 'OODA' loop following his experiences in the Korean War. So mission command provides a way of allowing operations to continue despite the Clausewitzian fog and friction of war in complex air campaigns with many assets in play, or where beyond line-of-sight communications are degraded or denied. It also – in theory – provides air commanders with the space to focus on the big issues of command without being 'fixed' by the need to control every detail of tactical execution. This explains why centralised control, decentralised execution, enabled by mission command, emerged after the Second World War as the favoured model for air C2.

The wisdom of this choice was apparently proved in the Cold War, when the freedom of action that western air forces enjoyed contrasted starkly with the Soviet predilection to centralise both control and execution. A rigid, Ground-Controlled Intercept (GCI) approach to air combat sapped Warsaw Pact pilots of initiative and was symptomatic of the inflexible, Soviet-bloc approach to air power that robbed it of much of its utility. Yet air C2 is always context-specific. In the 1967 Six-Day War, the success of the Israeli Air Force exemplified the triumph of the high-quality individual pilot in the cockpit exercising mission command over the system-based, centralised approach of his Soviet-trained and influenced Arab adversaries, who always had to wait to be told what to do. However, in the subsequent War of Attrition (1969-70), the *Heyl Ha'Avir* itself reverted to centralisation, with the Chief of the Air Force personally controlling execution from his central command bunker. This was because he had both the means to do so - the tiny size of Israel meant that all air operations took place within line-of-sight radio contact and the numbers engaged were small - and equally importantly, the motivation. Egyptian aircraft were often flown by Soviet advisors, so every tactical engagement decision implied potentially geo-strategic consequences. In these circumstances, the Israelis were simply unwilling to delegate execution when centralising it was both possible and, in this instance, desirable.

Why change?

While the Israelis could centralise execution during the War of Attrition, because of the limited range and scope of air operations, this was not a scenario that western nations could expect in an air war on NATO's central front in Europe. But since the end of the Cold War, the development of network enabled capabilities based on space-borne, beyond line-of-sight communication and data-transfer systems has created a new environment for air C2. This is a double-edged sword. On the one hand, it provides us with the ability to deepen mission command and decentralise execution if we choose to do so. We can use these capabilities to disseminate the commander's intent quickly and reliably down through the chain of command, using datalink to synthesise and project the best available information into the cockpit to support high quality decision-making at the lowest tactical level. On the other hand, this technology also makes centralised control more feasible, and the Israeli example demonstrates why this may be desirable (or even necessary) in some circumstances.

In any case, the RAF has long since departed from the simple purity of centralised control, decentralised execution, despite our stated doctrinal position. 'Operations other than war' have always attracted political scrutiny, and a generation of airmen who cut their teeth enforcing no-fly zones over Iraq in the nineties will be only too aware of the level of control exercised by the remote Combined Air Operations Centre, either directly, or through the airborne mission commander: 'Mad Dog directs'. And I would argue that we have never practised centralised control, decentralised execution and mission command in the way that we have preached; ATOs – like Bomber Command's op orders - that mandate package size and type, target and alternates, take off time, target and tanking brackets, vulnerability times, permitted weapon load-outs and flow direction do not, in reality, allow much freedom

of manoeuvre. So the publication of *JDP 0-30* in 2013 was timely, because it provided an opportunity to review air C2 from first principles and consider if there is still a single, ideal model now that technology apparently offers us the choice to centralise or decentralise as we see fit.

Towards a different Philosophy of C2

Perhaps unsurprisingly, the authors of *JDP 0-30* concluded there is no simple template that is applicable in all circumstances. As one size does not fit all, a new and more flexible approach is required instead, where centralised control is still relevant and necessary, but execution authority should be delegated to the point where the best level of understanding is available. Sometimes, this may be at the air headquarters or above, but on other occasions it will be more appropriate and effective to decentralise execution to the cockpit. This new philosophy of air C2 is characterised in *JDP 0-30* as *centralised control, adaptive execution*.

Adaptive execution acknowledges that a centralised approach may be more appropriate for small-scale missions where the stakes are particularly high or where the highest-value assets are in play; where the participation of unfamiliar or less-capable coalition partners, possibly with very different ethics and values to our own, makes it inappropriate (or unwise) to decentralise execution authority; and when a lower-level decision-maker may be unaware of the full strategic consequences of their actions. The latter may be a particular imperative when we are employing full-spectrum targeting and cross-component capabilities, as recent operational experience indicates that execution authority is unlikely to be delegated below the strategic level when we use a mix of lethal and non-lethal effects.

Any suggestion that we should centralise execution creates justifiable concern that we may sap initiative or erode agility, replicating the failed Soviet-bloc approach to air C2. But in reality, there is nothing new about centralised execution, so we should not necessarily shy away from it. Since the air attacks on the World Trade Centre on 11 September 2001, almost every western air force has developed mechanisms to centralise execution authority at the highest strategic level (usually political rather than military) if an air defence fighter is required to engage a hijacked civilian airliner, especially over densely populated urban centres. This is a good example of an air operation where the stakes are so high - in terms of loss of life and the political consequences - that centralised execution is not just desirable, but essential.

Any operations with a high political tariff are therefore likely to generate the 'long screwdriver effect' and encourage (or mandate) a centralised approach to execution. Consequently, air commanders must understand the operational implications, particularly in reducing tempo. In October 2001, the operators of a US remotely piloted air system pinpointed the supreme leader of the Taliban, Muhammad Omar, in a convoy of cars fleeing Kabul. As neither the operators nor the US Navy Fifth Fleet commander in Bahrain could authorise the strike, approval had to be sought from US Central Command in Tampa, but the ensuing delay allowed Mullah Omar to escape.

So in sum, within the framework of adaptive execution, decentralisation (enabled by mission command) remains the default ideal, because it is most likely to assure tempo and flexibility by shortening the decision cycle. However, adaptive execution also acknowledges the reality that we may have to centralise execution in some cases, even if this is not always desirable in purely operational terms.

So what?

In one sense, adaptive execution as a concept simply recognises the reality of current operational practice. In air operations over Libya in 2011, for example, authority to execute attacks was delegated up and down the chain of command on a case-by-case basis: to the combined joint task force headquarters or combined air operations centre; to joint tactical air controller teams on maritime patrol aircraft; or to the cockpits of strike aircraft. Effectively, this was centralised control, adaptive execution before the term had even been coined. However, Libya also highlights two particular challenges with adaptive C2.

First, the speed of communication is critical. When execution authority was centralised to the combined air operations centre, target locations were often passed as GPS coordinates or in verbal code, as secure communications were not available to all coalition aircraft - but this risked errors in transmission or misunderstandings because of language difficulties. So execution could be a lengthy process, especially if messages had to be relayed through third-party agencies, and this meant the ground picture could change before aircraft were ready to attack. In contrast, when execution was decentralised, we could exploit the 'eyes-on' advantage available in the cockpit to engage fleeting targets.

However, this leads on to the second challenge: determining where the best point of understanding is really located. On one occasion, the pilots of a flight of Typhoons were convinced they could positively identify pro-regime forces attacking civilians and requested engagement authority. However, the CAOC withheld execution authority, as it enjoyed access to credible, near real-time human intelligence indicating no pro-regime elements were in the area. Although the aircrew were disappointed that authority to engage was not decentralised to them, it subsequently transpired that anti-regime forces had mistakenly engaged each other in a 'green-on-green' incident, so in this case, the decision to withhold execution authority was absolutely correct.

These examples reinforce the point that no single model of air C2 is appropriate in every case and we must simply learn to live with ambiguity. This implies a cultural change, where we take the air estimate process more seriously and afford air C2 particular priority within it, so that we can select the most appropriate degree of delegation to suit the particular operational circumstances. We must also review delegations constantly, because in politically ambiguous or dynamic operations we will have to direct execution authority upwards or downwards for different phases of the campaign, particular missions, or even for different targets within the same sortie.

The bottom line

Air C2 is important. In an era of financial austerity, we cannot expect to improve our capability through a series of major platform procurements. Instead, we must fight smarter, and this includes developing a more sophisticated understanding of the processes and principles of C2. In the past, I believe airmen and airwomen have been guilty of not giving command the attention it deserves. The mantra of centralised control, decentralised execution as a 'one size fits all' approach contributed to this unfortunate state of affairs by encouraging a degree of intellectual laziness; it was easy not to think too hard about C2 when there was only one possible answer. As one of the Royal Air Force's senior leaders said after Operation ELLAMY:

> *"For the last twenty years our strategy has been to turn up and wait for the USAF to hand us the ATO. There was a tangible sense of shock when we realised this wasn't going to happen this time".*

This is not good enough in an uncertain future of potential contingent operations where we may have to think for ourselves. I believe we can use the concept of centralised control, adaptive execution as a vehicle to promote a more rigorous approach to the employment of air power that encourages air commanders (and their staff) to think harder about every air operation they undertake: what is the context, what are the operational circumstances, who is the opponent and what are the desired outcomes? Only then can the most effective model of air C2, with the most appropriate delegations, be determined to meet the objectives set by higher authority and deliver operational success. Clausewitz famously asserted:

> *"the first, the supreme, the most far-reaching judgment that the statesman and commander have to make… is to establish the kind of war on which they are embarking".*

As airmen and airwomen, we ignore this aphorism at our peril.

Notes

[1] DCDC, *Joint Doctrine Publication 0-30 'UK Air and Space Doctrine',* dated Jul 13.
[2] DCDC, *Joint Doctrine Publication 0-01 'British Defence Doctrine'* (4th Edition) dated Nov 11.
[3] Despite the numbers of aircraft engaged, in most circumstances only the formation leader spoke directly to the controller.
[4] *JDP 0-01, p 5-4.*