

# AERA 3 and the NAS Plan

## The Attempt to Automate ATC, 1981-1994

Gary G. Nelson (garygnelson@gmail.com)

June 11, 2024

### Abstract

This memoir records my participation at the Mitre Corporation in the benefit-cost (B-C) analysis of the Advanced Automation System (AAS) component of the National Airspace System (NAS) Plan and development of Automated En Route Air Traffic Control (AERA) phase 3 (AERA 3). My work on these programs at Mitre was between 1984 and 1993 before the AAS was canceled by Congress in 1994 despite favorable B-C. The NAS Plan for modernization of the air traffic control (ATC) system was originally published by the Federal Aviation Administration (FAA) in 1981 and regularly updated. The AAS would have provided an automation platform to update the NAS Stage A platform that became operational in 1973 and was based on IBM 360 mainframe processors. The AAS would have consolidated en route and terminal ATC radar data processing (RDP) and flight data processing (FDP) and enabled the sequence of AERA applications to improve ATC operations. AERA 3 would have automated the ATC level of what became codified (ca. 1990) as three-level Air Traffic Management (ATM). This experience also extended to the Automated Vehicle Highway System (AVHS) under the

Intelligent Transportation System (ITS) program in the 1990's. The development of AERA 3 and IVHS presaged many current issues in the automation of safety-critical systems (e.g., autonomous vehicles) but the truncation of Federally-sponsored development only delayed dealing with the risks.

## **A Golden Age of Cybernetics?**

The Semi-Automatic Ground Environment (SAGE, implemented in 1958) for air defense of the U.S. was the first large scale digitization and computer mediation of control and communications that became the Internet of Things (IoT). I was recently struck by how this legacy is being forgotten at a site that is now abandoned in the pine woods of the National Seashore on Cape Cod. There in 1956 was established the Experimental SAGE Subsector (ESS) that gave the Whirlwind Computer in Cambridge its sensory input (surveillance and height finding radar) and its means of actuating airborne interception by Ground-Air Transmit and Receive (GATR) linkage intended to automatically direct interceptor aircraft both human-piloted (the F-106) and automata (the BOMARC missile). While the Whirlwind site on the MIT campus has IEEE recognition, the ESS site that is the part making the prototype IoT lies recently demolished and unmarked. SAGE, NAS Stage A and AAS/AERA were all steps in the epochal development of cybernetic systems with various levels of human oversight. Cancellation of AAS/AERA was a fragmentation of ATM development but also a gap in coming to grips with challenges now posed by autonomous vehicles (AVs) and artificial intelligence (AI) generally.

The SAGE system was developed by a chain of institutions from MIT to MIT-Lincoln Labs to the MITRE Corporation formed in 1958 for SAGE system integration. Mitre subsequently adapted the SAGE technology to automation of en route ATC as the National Airspace System (NAS) Stage A

implemented in 1973. Terminal RDP systems (they do not do FDP) were separately developed, originally analog and at levels of capability in the automated radar terminal systems (ARTS). My own career ended in the terminal automation department of FAA in 2011.

The technology steps from computer-mediated communications in SAGE to ARPAnet and then the Internet are well known. They are classic examples of the government relation to technology development in defense applications. The civil system in NAS Stage A has had an increasing air defense function as the early airborne warning and control facilities of the Air Force were consolidated. Personnel and expertise from that system were still active in my FAA department when I was there 2007-11. Both air defense and ATM are unique Federal functions relative to private commerce and air navigation. The step from SAGE to Internet however bypasses the question of human participation in safety-critical systems. Use of the Internet and its “social media” applications have just diffused the issue of accountability for risk both in the messaging and security of the channels. My experience, that also included a stint in the Homeland Security Institute (HSI, the original version 2004-2007) causes me to reminisce about the Federal role in technology and the distribution of accountability for risk as technology extends our interactions. We are living in many extensions of the concepts originally called “cybernetic” [from Wiener's book, 1948] but the original concerns around cybernetics [including per Wiener's title, “The Human use of Human Beings”, 1950] have fragmented accountability. AERA 3 development on my part focused on the scale-hierarchical allocation of accountability for risk, with a place for automation (we would now say AI) and a place for human beings. [I am currently working on a text concerning accountability for risk and equity in the governance hierarchy, with particular attention to transport projects].

# System Engineering in Complexity

After the cancellation of the AAS part of the NAS Plan in 1994 I left Mitre only to return in 1996 on the ITS work and then a varied career ending in the FAA. In that interval I was aware of the continuing debate about what happened to the NAS Plan that had been a deliberate program for eight years.

As disclaimer, i have discussed the impressions here of what happened among other involved people and there is no clear consensus on “what went wrong”. My perspective and conclusions from the B-C analysis, over two phases and as liaison to the (then) Congressional Budget Office (CBO, now GAO) and further study of complex systems may be unique.

The NAS Plan was initiated in 1981 by FAA Administrator Lynn Helms although the first “public” release was dated to 1982 [A Brief History of the FAA at [https://www.faa.gov/about/history/brief\\_history](https://www.faa.gov/about/history/brief_history)]. The NAS Plan went through annual updates including a Congressionally-mandated update of 1985. The last [1989] NAS Plan can be found online [at <https://apps.dtic.mil/sti/citations/ADA215882>]. As the FAA states [op. cit.]: “In February 1991, FAA replaced the NAS Plan with the Capital Investment Plan (CIP). The name change subtly reflects de-conflation of NAS Plan components that in turn reflect on the process of managing a complex system, indeed an “ecology” to use a more “modern” term that would also be anachronistic then. I once had a complete set of the annual NAS Plans left behind when I left the NAS work at Mitre in 1995. I have in hand the 1985 version [National Airspace System Plan: Facilities, Equipment and Associated Development, USDOT-FAA, April 1985, Congressionally Requested Update]. As that version states [pg. 1-3]:

“FAA believes that this Plan will achieve dramatic service improvements, represent[s] a practical way to achieve a significantly safer and more efficient system and will reduce FAA costs to the taxpayer.”

I joined Mitre (Civil Systems Division, McLean, VA) in 1984. Mitre is a collection of Federally Funded Research and Development Centers (FFRDC) for various Federal patrons. The ATC work had been a spinoff from the SAGE Air Force work but the FAA work gained its own FFRDC as the Center for Advanced Aviation Systems Development (CAASD) in 1987 while I was there. I was hired onto the B-C analysis for the AAS, a cross-discipline task of economic analysis and system engineering for which I had experience. I focused on operations and maintenance (O&M) effects and later on the Area Control Facility (ACF) effects. The first phase of the analysis was aimed at implementation of the Host computer system, the replacement for the 9020 (IBM 360) mainframes that was accomplished by 1989. The second phase of the B-C analysis I worked on 1987-89 was re-baselined to Host operation.

Diagrams of NAS Stage A [figures appended] show that despite mainframe platforms there was nothing monolithic in the hardware or functional architecture. The reliability and information bandwidth requirements already implied multiplicity of system modules all connected by a multiplexed channel of the Peripheral Adapter Module (PAM). RDP processing is a separate thread. Originally outage of RDP was addressed by placing the sector displays horizontal with plastic “shrimp boats” used to identify analog radar targets from the flight strip data. That was pre-NAS Stage A continuity now as backup from the time when en route ATC was all flight-strip based. NAS Stage A both digitized the RDP and with the FDP thread correlated flight ID with RDP tracks. What was “wrong” there is the persistence of paper flight strips. The ability to go beyond sector scope of control with FDP is what AERA would tackle.

The results of the B-C analysis may help interpret the Congressional decision on AAS but much more was at stake regarding the integrated program articulated in the NAS Plan. The main effect of AAS cancellation, and so NAS Plan disruption was to fragment and further delay the proposed changes. After 2000 a more comprehensive scheme for the air transport system (the whole aviation ecology) became known as NextGen (Next Generation Air Transport System), Integration and a comprehensive architecture were good. The problem was the distributed responsibility for actual projects. The ATM part under the FAA continued in pieces for separate en route and terminal systems and the communications/surveillance components. Only when I left the FAA in 2011 was the ACF idea resurrected. In the meantime the semi-analog ARTS was being replaced incrementally by the digital (and distributed-processor) Standard Terminal Automation Replacement System (STARS). I was in the program office that included STARS and only by the time I left was the system going to re[place the ARTS III installations at the largest terminals. Meanwhile the en route automation renovation was ongoing (En Route Automation Modernization System, ERAMS). That is nearly two decades after the intended NAS Plan schedule.

There is a sidelight to the fate of old NAS Stage A. Withholding STARS from the larger terminals was a risk-shy strategy that extended ARTS tenure that went back as far as the NAS Stage A [ca. 1973]. The smaller ARTS I and II systems were analog until replaced by STARS. They basically overlaid target-tagged radar images on a TV projection of the airspace map. But the NAS Stage A started with two versions of en route RDP. The “standard” version used a “high speed switch”, a special processor for the digital radar displays. However, as an alternative some facilities had a version of the 9020 (IBM 360) processors for the RDP. As a result when the display channels were being replaced in the later 1990's some of the IBM 360 hardware was still in the system. While “old technology” had become

a cliché for FAA before and after the AAS initiative, in my opinion the persistence of ARTS III and the display channel IBM 360 technology are the worst examples. The ATC system was basically extending 30-year old computer platforms. As the B-C analysis showed the issue was less about hardware wear (aging) than functionality to air transport.

The NAS Plan was to address the tendency to keep old technology and their limited applications. It was hardly avant garde in that respect precisely because any system transition would be long and involved critical operational risk. The Host was only a hardware replacement for the same function. I add that one of my last tasks (ca. 2010) was integration of Automated Dependent Surveillance Broadcast (ADS-B) locational messages into the terminal RDP. That was a good tracking of the technology capabilities as part of NextGen and ADS-B also enables the public apps for flight tracking on your personal platforms. However the special role of ATC was at issue in some proposals to replace the radar surveillance system by ADS-B rather than “fuse” the two target location methods. An ADS-B only system would have ignored the defense function (as do the perpetual calls to privatize ATC).

ADS-B would also supplant the last project I was a manager of: A special feed from airport surveillance radar (ASR) to some airlines that used the locational information on their flights for more efficient ground management. I also made a failed attempt to get FDP into the terminal automation, in reaction to what was still manual fetch of paper flight plan information. It was my conviction that improvement would be necessary for at least one application being proposed by Mitre called Relative Position Indicator (RPI, but not my alma mater). About that time I also worked on some “open system” protocol standards for ATC information.

These late technology applications and the “open systems” concept reflect back in time to the collision

of ATC with the rate of technology change and return to the issue of what happened to the NAS Plan and why AAS and AERA was abandoned as an integrated program. That is about how the concept of “complex systems” and their “evolution” was starting to exercise the profession by the late 1980's and so the time when the PC and Internet were penetrating even the general public. Mitre had itself seen the changes. When I arrived in 1984, the ARPAnet could be accessed at our node by using a video terminal in the hall hooked to our corporate server (a small mainframe computer). One could follow a newsgroup via ARPAnet. That did not much interest me at the time and only in the ITS work (ca. 1996) did I latch onto the alt. classification hierarchy of topical discussions on the Internet regarding transport-network modeling. Shortly after 1984 Mitre installed desktop terminals to the corporate computer and only by 1987 did I have a PC for a research project. I had had an Apple II+ at home since 1982. Clearly the technology and systems-concept winds were blowing against the still-mainframe organization of the ATC system. That would have changed with the AAS but while RDP processing could be distributed FDP was inherently a centralized function but with distributed input/outputs..

At the time I became interested in what I will loosely call “complexity science”. That started on the research project that was about prediction of flight progress with the winds aloft and that, surprisingly, leads to my part in AERA 3. To summarize, the winds problem was prediction over space-time scales and like general weather prediction involved the scaling from local state measurement (observations) to global observation assimilation and prediction via atmospheric models. That was my first exposure to a scale hierarchy (i.e., the Kolmogorov power spectrum of what is approximately homogeneous turbulence). In subsequent work, including AERA 3, ITS and homeland security I developed the scale hierarchy in application to risk management. The ideas can be glossed in two of my papers [Gary G. Nelson, Organizational Evolution, Life-Cycle Program Design, Washington Academy of Sciences,



Volume 92, Number 1, Spring 2006; The ConOps in a Self-Similar Scale Hierarchy for Systems Engineering, 2007].

The basic ideas of a scale hierarchy I borrow from Stan Salthe [Stanley N. Salthe, *Evolving Hierarchical Systems: Their Structure and Representation*, Columbia University Press, 1985; *Development and Evolution: Complexity and Change in Biology*, MIT/Bradford Book, 1993; *Hierarchical Structures*, *Axiomathes*, Volume 22, Number 3, September 2012, pp. 355-383; *Perspectives on Natural Philosophy*, *Philosophies* 2018, 3, 23; doi:10.3390/philosophies3030023]. I have had many exchanges with Stan that meet his ideas from natural systems with mine in physical science and engineering. At the time I was understanding Salthe (ca. 1989-1992) I was also involved in the Mitre Nonlinear Systems interest group and the Washington Evolutionary Systems Society (WESS) that collected international experts on the topics. This is to say that I was getting a perspective on some of the concerns that were also assailing the NAS Plan's approach to a "complex system" in the "turbulence" of rapidly developing (or "evolving"?) technology.

The AAS had after 1989 already gone through a design competition phase (that I also participated in) and subsequent prototype development contract, going to the IBM Federal Systems team (relict of the SAGE and NAS Stage A work), shortly taken over by Loral and subsequently to what became Lockheed-Martin that still had the en route system development when I left in 2011. The AERA phases continued under development at Mitre. Yet after the cancellation of AAS and the morphing of the NAS Plan to the CIP the accusation levied on it by many of my colleagues was that "it" (with some ambiguity between the entire NAS Plan and its AAS part) was too much a "big bang" program in contrast to the emerging idea of evolving systems. It is that accusation I examine as a case of the NAS Plan colliding with new and still immature ideas about complexity and project management.

## To Evolve or to Let Evolve

I start by testifying that the term “architecture” had not been used regarding the ATC work at Mitre prior to about 1988. I first heard the term from colleagues that came over from the part of Mitre (next door) working in the Army command and control FFRDC on the Worldwide Military Command and Control System, (WWMCCS). Before that we would have said “system level specification” (SLS) in the semiotics of system engineering since its codification by the Air Force after the missile and SAGE programs. These military standards (mil-std) were also morphing to the “commercial” standards for system engineering under the IEEE and further into the management process “quality” standards of ISO. The IEEE defines architecture [IEEE, Std 1472, 2000, Recommended Practice for Architectural Description of Software-Intensive Systems, October 2000] as “The fundamental organization of a system embodied in its components...and the principles guiding its design and evolution.” The military was continuing with its own DoD Architecture Framework (DoDAF) as part of its system development process, the Joint Capabilities Integration and Development System (JCIDS). The JCIDS was a dominant approach in our HSI work after 2004 and DoDAF was widely adopted including for the NextGen architecture. These developments represent the transient in “systems thinking” with which the NAS Plan collided. But what was the substantive change and how much did that really borrow from the “science of complexity”? My ConOps article argued that the elaborated systems methods missed the essential point of complex systems of which the NAS Plan was no more or less guilty.

The “complexity concepts” affecting the systems engineering community still have to be parsed carefully and enter the arguments we had about the NAS Plan and the AAS. There are two aspects I focus on: The biological metaphors of development and evolution as widely misapplied (then if not

now) and the affinity to Herbert Simon's "The Architecture of Complexity" [1962] as articulated in the architectures of open systems, the Internet and WWMCCS. Simon's concept becomes the scale hierarchy as articulated by Salthe.

I use Salthe's scale hierarchy as pairwise levels of interacting modules (the focal or polycentric level) and a formal level. The formal level will be taken as the "architecture" with its protocols of interaction common to the focal modules. This idea is nothing more than the "network" concept defined in graph theory as the formal  $N$  and the focal bipartite set  $N = \{\{E\}, \{V\}\}$  where the  $\{V\}$  set (nodes or vertices) are the modules and the  $\{E\}$  set (edges or links) the interactions. This is how we structure systems and is the practical definition of complexity: There are many interacting parts that emerge as a "whole". The difficulty is in the dynamics of how the formal level (a new object) emerges from prior parts and then matures into stable form and common protocols of the  $\{E\}$  set. Simon only alluded metaphorically to the process (the blind watchmaker), called diachronic by Salthe and still under study in "complexity science" and for self-organizing (adaptive, hierarchical) systems that are as well called ecologies [as I owe to Atsushi Akeru, *Calculating a Natural World, Scientists, Engineers, and Computers During the Rise of U.S. Cold War Research*, MIT Press, 2008].

With this understanding of "architecture" the actual work of projects gets embroiled in the dynamics of complexity that summon the biological metaphors of development, evolution, and then their conjunction in evolutionary development (evo-devo). Biology makes a clear distinction:

- Genomes *evolve* as the form of their instantiations (modules as organisms).
- The organisms (phenotypes) *develop* from birth to maturity and then interact at the focal level.
- Evo-devo refers to the form of the process by which the genome is interpreted into construction

of the phenotype by proteins (from genome to proteome).

These concepts were all available (evo-devo being most recent) at the time we were disputing just what happened in the collision of the NAS Plan and the AAS with the “new” systems engineering. It is somewhat anachronistic for me to summarize the concepts I later wrote about to characterize the debate about the “big bang” at the time but they help to clarify the issues discussed among us engineers. They do not characterize the action of Congress, for which I refer more to the B-C analysis and what I learned about the process at that level from some 27 years in the DC orbit including the ITS and HSI projects and so around DHS, USDOT, FAA and FHWA headquarters.

I start with my objection to how “evolve” was being used in contrast to “big bang” regarding the NAS Plan and AAS. Some of my colleagues started to talk about “evolving a system”. Regarding the biological metaphor using “evolve” as a transitive verb is a fallacy. It is an important fallacy that implies evolution as an efficient act by an agent, also the fallacy of teleology (that a formal level has an intent that explains the prior focal modules). The term “development” was avoided, but that is just what every module does and what every task of a project is about. We do not “evolve” systems but complex systems do evolve and their modules develop. Evo-devo could be applied to the system engineering doctrine itself, but that was indeed formal to the focal debate about what happened to the NAS Plan and the AAS that I recall here.

An “architecture” then is a codification of modules, in a functional specification hierarchy (e.g., as in the evo-devo of object-oriented programming). That is consistent with the later IEEE definition. I later worked in the ITS architecture and affirm that as for ATM or NextGen the functional modules had previously evolved and continued to develop with or without the codification of an architecture. But the role of the formal level—both architecture and protocols—in open systems generally is important

for allowing any *development* to parse a project (concerning a set of modules at some specification level) within protocols that deal with the larger system. This assumes a quasi-static form (i.e., the focal development does not alter the external system). There is however the complication of being in a transient when the formal level and protocols are still emerging. I was closely embroiled in that in the ITS and HSI work (where I helped define a new function for the ITS architecture and managed a specification hierarchy of just what the homeland security functions were). I do not think that issue of a transient was essential to the NAS Plan case except when we came to AERA 3 and the need to parse human and automated functions as I will recall below.

This summary of what came of the “new systems engineering” is to argue that it had no substantive effect on the fate of the NAS Plan or AAS. It is not a “big bang” to coordinate several module developments under the NAS Plan still operating with the old system level specification (SLS) concept. The AAS was not a “big bang” development but had the challenges of system integration that any project has for a “complex system” (i.e., ecology or parts in a whole). How to parse and manage projects was and is the challenge. My own impression (as of 2011) was that NextGen was too much architecture codification and too little development coordination. The fragmentation of the NAS projects after 1994 to what last I knew in 2011 is a symptom of the failure in management from Congress on down of the project(s) implied by the NAS Plan. The failure in continuity of ATM applications development with the interruption in AERA is a failure of program management that has echoed in the recent challenges of risk management for AVs and AI. To communicate (however open the Internet and other open systems) is not enough. There are functions of real and risk-critical actions that projects and systems have to manage.

My experience, in several projects and functional domains, leads me to conclude that the lessons of

complexity have not been adequately applied to a scale hierarchy that extends from our Federal governance on down. The defects in the NAS Plan and AAS were not in their scope nor the fact that the ATC system had not yet become an “open system”. The defect was that the management of that program was caught in a transient of doubt about systems subject to rapid technical changes offering novel development options. The ATC system was highly conservative in its structure, technology and protocols. It was an open system by definition of having well-defined and discrete modules and codified protocols of interaction. That the protocols were variously hard-wired and not open to outside systems was a matter of history and security. That was going to change anyway. It was just a question of how fast the developments under the NAS Plan could occur and that was always a tradeoff of operational and project risk in the dimensions of cost, schedule and performance. Those are the issues discussed with the B-C analysis.

## **Benefits and Costs to Whom?**

I happened to get into the B-C analysis of the AAS through my system engineering training and work on a B-C manual for rail projects. I knew little about ATC when I arrived at Mitre in 1984. I rapidly had to learn about some 120 separate modules of the ATC system that could have been affected by or affect the AAS and O&M costs. That applied as well to the analysis of the ACF consolidation that affected all communications links.

There were two phases to the B-C work. The first we completed in mid 1985 and was used as part of the Host acquisition. Then a hiatus of about a year when I worked in the ATC communications division, mostly on the use of the analog radar channels for data communications, called the Radar

Microwave Link (RML) in the NAS Plan. In 1987-89 we did the second phase B-C analysis and toward the end I was liaison to the Congressional Budget Office on the results. After work on AERA 3 (1990-93) I saw the end of the AAS in sight and had left that program by its cancellation in 1994. I could understand the risk and policy issues while not agreeing with the Congressional decision, but neither I nor the B-C results were eligible to question Congress. Mitre in any case was only client to the FAA patron however confused that relation was.

Mitre carried the technological flame of SAGE to the FAA and therefore held strong systems-knowledge against the operational knowledge and risk accountability in ATC. By the time of the NAS Plan, about a decade after NAS Stage A development, many of the Mitre engineers from that generation were still around and held the domain knowledge. The AAS initiative was about as bold as the original NAS Stage A and about as much a change in the “architecture” of the ATC system as NAS Stage A was from the prior analog RDP and teletype FDP. NAS Stage A was still an en route-only ATC system and used paper flight strips although with automatic sector (i.e., radar display) distribution. Mitre was in a position to continue its “integration” role if FAA committed to the program. The issues of the benefits of the AAS against cost were similar to those of NAS Stage A that rather rode the technological wave from SAGE. I cannot first-person attest to how much anguish there was about benefits, cost and project risks for NAS Stage A that evidently was committed to. I do know that the engineers of that time were fully involved with the contractors (IBM and Raytheon) in the system development. IBM held the experience of developing Whirlwind into the SAGE system and so to the 9020 system of NAS Stage A. Raytheon had come from a different development path to development of the RDP channels (excepting the alternative based on IBM's 9020 processors).

It is relevant that by the time of the AAS, after 1985, there were two policies that kept Mitre more at

arm's length from AAS development. One was the System Engineering Integration (SEI) contract that came into effect one year after I came to Mitre. That went to an umbrella contractor that had to scramble to collect domain experts and started a trend still going on by the time I left the FAA in 2011. The Federal government after the 1970's (and two epochs of gas and economic crises) was in a "mood" of budget constraint. I say "mood" to hide a complex of real and ideological constraints I will not elaborate. It is hard to compare that "mood" with the clear budget conservatism of the Eisenhower years that nonetheless gave us the ICBM, SAGE spy satellite and other technology-developing programs. But anything that Mitre and the FAA did in the 1980's was now in a new governance and budgeting context. Within that was the long term trend to convert Federal-employee and FFRDC budgets (and expertise) to frangible contracts and contractors. The SEI was the initial version of that for the FAA. The entire contracting trend represented a fragmentation of program management and technical expertise. That is not to denigrate the role of implementation-contractor expertise as applied in SAGE and NAS Stage A or ARPAnet. The issue is project management and technology integration. The "golden age" of cybernetics was over. There was a real change in national economics and innovation begging clarification of roles in the scale hierarchy of governance and management. It is the same problem of accountability in human-automation systems.

The immediate effect was that there was much closer budget scrutiny on the NAS Plan and the AAS while the prior "integration" role handed from MIT-Lincoln Labs to Mitre was also hobbled. AAS went to the design competition phase between two candidates (led by IBM and Hughes). Mitre could only participate in proposal evaluation. There was a widespread impression that parts of both proposals should have been integrated but Mitre could not perform the "hands dirty" role that had occurred for both SAGE and NAS Stage A. Later when I was in the FAA I had a hard enough time overcoming the contractual stovepipes between the FDP in the en route system and requirements in the terminal



automation system. It remained my opinion that some FDP had to be integrated into the terminal ATC and after all that was implicit in the ACF concept of the NAS Plan and its renaissance after 2011.

The B-C analysis was an opportunity to “nudge” system design according to cost and benefit (via performance) evaluation. That was an integration of engineering and economic analysis I was qualified for. I add that after our work, Mitre also split any economic analysis off into another division losing that kind of integration. But then it did not work for the AAS either.

The primary issue for the B-C analysis then came down to that point cited in the 1985 NAS Plan: “...and will reduce FAA costs to the taxpayer.” This is a significant point of how distribution of benefits and costs is handled in B-C analysis. The doctrine is (and I have other disagreements with it) that the B-C aggregates all effects into pseudo-dollar totals. Who actually bears cost or gets a benefit is a separate issue. The FAA statement and Congress' budget interest is however on the internal effect to the ATC operation: Would modernization and development save real budget money? And the answer we found to that was: No. I do not think the answer was different for SAGE or NAS Stage A. But the difference in 1994 reveals much about what had changed in the nation. Our B-C analysis had several components, somewhat changed between its two phases:

- The capital cost of system acquisition
- The O&M cost (excepting controller staffing)
- Reliability benefits to system services
- Controller staff costs
- Airspace user benefits (equipment costs under the acquisition and O&M costs)

These components were divided into our analysis work tasks. I had the O&M costs and the equipment

costs for the airborne systems (starting with the transponder upgrades). Others worked on the other components that we combined in what was initially a large spreadsheet (with hard copy taped over walls) and later under my innovation of custom Fortran formatting. I add that was from my experience that came from the era of formatting word processors by line codes and for line printers. In the second phase I also had responsibility for the ACF impacts involving facility and communications costs.

The acquisition costs underwent an important change between the phases. In Phase I we used the FAA budget estimates. In Phase II we performed an independent risk analysis on costs. I no longer have the exact results. However the expected cost we predicted in 1987 was substantially higher than the FAA estimate and I recall looking at an Aviation Week article ca. 1994 that reported the “inflating” cost (on the order of \$3 billion) on which the program was ostensibly canceled. I remember that the updated budget costs were very close to our risk prediction from seven years before. It was a failure of risk management.

My prediction of the O&M costs even in Phase I was the first puncture into the NAS Plan goal of reducing internal costs. It was anticipated that surely modern equipment would have reduced O&M costs. In part that was based on a fear of wear failure in the 9020 processors. Remember they were IBM 360 mainframes with card-deck input. The 360's were built from encapsulated modules of discrete components, not “chips”. There was a concern of silver migration from the silver solder that would short the circuits, hence increasing failure. Perhaps. But there was no proof and no data to confirm wear effects on computers. The AAS was configured with many contemporary non-mainframe processors to improve performance and availability over the dual-thread 9020 system (perhaps). However the main cost effect was from the display configuration. NAS Stage A had one cathode ray tube (CRT) display per sector with small auxiliary display, the flight strip printer and a board to array

the strips. That was a few analog-electronic things to go wrong. The AAS called for a sector suite of trifold and larger displays with local processors for RDP and FDP interface. The tradeoff in O&M costs for more “glass” and processors was not self-evident but my analysis concluded that O&M costs would increase, not decrease (on a life cycle basis against the NAS Stage A baseline and especially with Host replacement of all 360 processors excepting those display channel installations that in any case were with us another decade).

Regarding controllers and other costs, the effects were entangled in the sector configurations and the ACF consolidation. Clearly the ACF would increase communications by the high-reliability (multi-path) homing of all communications to less-local ATC facilities. The 180 or so terminal controls would go into 23 or so ACFs, roughly the number of existing en route centers. An option I favored was to keep a set of large terminal facilities as backup. Then many terminal facilities and equipment were eliminated against increased ACF costs. But an important factor was that controller salaries were based on facility traffic: The per-staff cost was going to increase against possible staff reductions from consolidation. Current consolidation plans should be consulted but for the AAS the net change in staff and O&M costs internal to the FAA budget were at best neutral compared to the large capital (and risk liable) cost.

The airspace user benefits were from the AERA implementations that would follow the AAS platform. Our prediction was that these were positive enough to well outweigh costs (all with proper time discounting) even with the risk-expected cost. The user benefits were not my task although I did manage a survey to analyze current flight deviations from optimal paths. AERA was intended to remove any ATC-imposed constraints on flight-optimal paths. Given traffic growth even small savings added up to big benefits. In my opinion the costs from traffic flow management (TFM) constraints

would be harder to reduce but that concerns my AERA 3 work and my later work on both TFM and airspace design that in the Airspace Management (ASM) function forms the three levels of ATM as defined by ICAO:  $ATM = ATC + TFM + ASM$ .

Several of us were outraged by the Congressional cancellation of the AAS, apparently based on increasing budget costs despite large public benefit. That is a matter of how you count distributed costs and benefits, and that is an OMB and B-C doctrine matter beyond my authority. Whatever your opinion, the death of AAS fragmented further modernization and functional improvement of ATC and, as I argue, improvement of the entire ATC-TFM constraint on air traffic that NextGen is so concerned with. It is a serious matter of how our operation of the air transport mode is managed and how technology is applied to it. The AAS was a failure of that management, probably increased overall costs of the fragmented programs and delayed technology applications.

In 1993 after leaving the ATC work at CAASD I worked for a year on program risk for the Army's chemical demilitarization program, also at Mitre. The chemical processing industry in particular had tracked and modeled growth in project schedule and costs and reductions in performance over time from initial predictions to completion. Not surprisingly there is an approximate logistic (S-shaped) curve from initial predictions to the realized values. Our cost-risk analysis for the AAS just exposed this trend for AAS costs that should have been recognized by the FAA, Congress, the GAO and the OMB. The cost "escalation" by 1994 and schedule delays should have been well anticipated in program risk management. Any surprises represent a failure in program management. And by the way, the disposal of the chemical weapons stockpile was supposed (by treaty) to have been completed by year 2000 when we were analyzing it. It took until 2023. The performance risk of AAS was never tested but the expectations bring us to AERA.

## AERA 3 and System Complexity

As stated in the 1985 NAS Plan [En Route System, pg. III-6]:

“The first implementation of the advanced automation system computers will occur in late 1992 with automated en route air traffic control (AERA 1) software operational in 1993...By the year 2000 area control facilities will be operational systemwide and AERA 2 and 3 will have been implemented sequentially.”

The development of the AERA phases was well underway when I arrived at Mitre in 1984. As I became involved in AERA 3 after 1989 it was doubtful that the 1985 schedule prediction would have been met, but that should have been no surprise. In the event very little of AERA got onto the operational en route ATC. The only AERA function I knew to be implemented by 2011 was a version of the Flight Plan Probe. That was one realization of the AERA intent to bring more of FDP into the sector ATC enabling clearance of adaptive and efficient flight plans by lookaheads and negotiation over the necessary string of sectors. AERA was never intended for the inherently local terminal ATC but the RPI application I helped coordinate among facilities was for better approach sequencing of flights by trajectory projection and display (called “ghosting” when a predicted target location is displayed).

It should be apparent that the trans-sector scale of AERA merged with the TFM scale. My work on AERA 3 was focused on outreach to controllers in a ConOps process and conceptualization of how TFM would interact with automated ATC. It can be appreciated that automating ATC (an end to human sector controllers) was a sensitive risk and employment issue for the human controllers. AERA had a set of AERA Controller Teams (ACT) and so ACT-3 for AERA 3. I think we had some success in

convincing the operators that it was advantageous to take humans out of the high bandwidth (or “high pucker” in the ATC vernacular) role of tactical flight control and into a supervisory role over the automation. This is of course the basis of practice in Supervisory Control and Data Acquisition (SCADA) systems as used for power systems etc.

AERA 3 would have been the late realization of the potential that was in the initial digitization of RDP combined with FDP as achieved in NAS Stage A. My colleagues at Mitre who were around from that time were generally cognizant of what they considered lost opportunities to exploit the digitization in practice although the AAS was necessary for much more. One of my colleagues present in the NAS Stage A work showed me a 1973 report that presaged AERA 3. That and other conversations showed that the issues raised by SAGE and NAS Stage A were well appreciated at the time of those pioneering systems. As one colleague stated (and I misquote from memory): "We gave air traffic controllers the automation but what were they going to do with it?" The clean, digitized version of an analog radar scope was much appreciated by controllers to make the existing machine-human interface better. The phases of AERA (nominally a basic AERA then AERA 1, 2 and 3) were intended to *gradually* exploit a re-allocation from human to automated decisions. What had been automated in NAS Stage A was referred to as “snitch patch” or properly Conflict Resolution Advisory (CRA) that automated detection of flight track conflicts if the controllers had not corrected them. The mixed human reaction to that is relevant to the AERA 3 concept.

By way of operational risk recall the ATC-mediated airborne collision over Überlingen Germany [July 1, 2002]. The apparent cause was human-controller override of a collision avoidance directive from the airborne Traffic Alert and Collision Avoidance System (TCAS). The development of TCAS played a large role in the AERA 3 concept and later the NextGen concept derived from what in the 1980's was

called Free Flight. It was supposed that tactical safety could be allocated to aircraft with adequate communications and automation. In retrospect that was a powerful motivation for AERA 3 in a supervisory and so combined ATC-TFM role. The later ADS-B development (that I was involved with ca. 2010) added capability to the concept although the vulnerability of a GPS-dependent system has to be considered. Both failure risk and the issues of the human role have to be considered in any ATM evo-devo. These issues of course re-appear for terrestrial AVs and the use of uninhabited aerial vehicles (UAVs or drones). AERA 3 work started to address these issues and its termination if to be regretted beyond the fate of the AAS.

I estimate that about 150 people were involved in the AERA development at Mitre ca. 1990. The Mitre staff were heavily involved in software specification and prototyping. For AERA 3 at least the software building was by contracted staff I include in the total. I think that the arms-length between we engineers and the coders was an unfortunate generational gap that impeded development. That changed with younger staff at Mitre as I could observe by 2011. I will admit that on my part (mostly a FORTRAN programmer of ca. 1970 on that IBM 360 system) it was not in my capability to translate my ideas about AERA 3 directly to prototype demonstration. However it is the cancelation in 1994 that interrupted any further demonstration of the concepts.

Regarding the ATM scale hierarchy, one of my last projects at Mitre concerned airspace design at the ASM level. Later in contract work (ca. 2002-04) I worked directly with the TFM office of FAA and so gained much more familiarity with that function. That was my second and unsuccessful effort to promote the AERA 3 concept of TFM-ATC inter-level coordination. However I was impressed with how manual the TFM function still was. There is the ATC System Command Center (ATCSCC) that has national flight oversight and was essential in the response to 9/11. We have the ground holds on

flights that replaced the "holding patterns" of flights at the airspace bottlenecks that are primarily at the destination terminals. That and the residual "airways" structure of the airspace (in large part reflecting the ground ATC sectors) is what AERA and then NextGen are reacting to in favor of flight efficiency and so the user benefits predicted in the B-C analysis of AAS. I see AAS/AERA and NextGen in a continuity that was however broken and then fragmented.

The basic aims of AERA generally to extend the scope of sectors and so to span the ATC and TFM scales had to face complexity in the airspace system and its "ground" information network. That is why I use the term "ecology" now although anachronistic to my work at the time. The ATM system is literally an eco-logy, a logical concept of what was happening in physical airspace and of "moving aluminum". It is its own part of the airspace and transport ecology. The transient of views on this ecology that AAS had to weather mirrored complexity as an intensification and extension of all interactions in the ecology. That is what motivated flights (incorrectly I think) to divorce from ATM while AERA was arguing for surpassing the ATC sector modules. The ACF concept was also implicitly arguing for less partition between en route and terminal ATC. The historical partition of ARTS and NAS Stage A development prevented full exploitation of that issue at the time. Indeed up to my retirement in 2011 en route and terminal remained different domains in system and FAA organization even if the automation capability was converging after STARS (initially ca. 2005 but still deploying during my tenure).

My tutelage in complex systems and the scale hierarchy defined how I jumped into the AERA 3 development about midway when I entered in 1989. I was not going to get embroiled in software specification or the ongoing prototyping of what would appear on the sector suite displays. I was instrumental in promoting two concepts within the Mitre group and to ACT-3:



- The complexity required a supervisory and not tactical-control role for the humans, and;
- The inter-scale interactions between TFM and ATC would also have to be automated once there was no human at the ATC-sector end.

These were, in my view, the same problem and would not be addressed by more ATC code writing. Consequently there was no depth of prototyping available to develop the ideas. They were implied by but outside the development paradigm of AERA. I alone lacked the capability to show in effective simulation what was needed and so ended up a decade later (unsuccessfully) arguing for the necessary development to the TFM part of the FAA.

At the time and since I was following the model of hierarchical regulation (or adaptive control). That was alluded to in cybernetics and developed in articles and texts I had found [George G. Lendaris, On the Definition of Self-Organizing Systems, Proceedings of the IEEE, March, 1964; Mihajlo Mesarovic, Self-organizing Control Systems, IEEE Trans. on Applications and Industry, 83(74), 265-9, September, 1964; Mesarovic, Macko and Takahara then published a text [Theory of Hierarchical, Multilevel Systems, Academic Press, 1970; George N. Saridis, Self-Organizing Control of Stochastic Systems, Marcel Dekker, Inc., 1977; W. Findeisen et. al., Control and Coordination in Hierarchical Systems, John Wiley and Sons, 1980]. Lendaris [1964] cites 1962 works on self-organizing systems by Ashby and Mesarovic, as well as by Zadeh on adaptivity [L.A. Zadeh, "On the definition of adaptivity," Proceedings IEEE (Correspondence), vol 51, p.469; March, 1963; W. Ross Ashby, "Principles of the self-organizing system," in Principles of Self-Organization, Pergamon Press, New York, N.Y., pp. 255-278; 1962; Mihajlo Mesarovic, "On self-organizational systems," op. cit. 1962].

The cybernetic-derived idea of multi-level systems was then recognized but for me it was Salthe's

general scale-hierarchy model that extended the ideas to ecologies including of human governance and “natural” participants such as the winds as they affected flight progress. At the time however not even TFM was a mature operation. It was generally recognized that TFM became distinct from en route control with the controller's strike of 1981 when indeed the sector sets had to be supervised at a higher level with the bandwidth limitations accompanying loss of so much sector-level control. That was a reversion to the older FDP-based en route ATC. At the ATC level the primary object was the flight track (the digital extension of the RDP target) within the scope (in two senses) of the sector. At TFM and FDP level the object is a flight, with waypoints, and then the aggregate of flight traffic compared to airspace capacity (partly determined by human sector workload). However it was only ca. 1990 that ICAO even defined the three ATM levels formally. My view is that there was a theory of hierarchical systems available but the institutional recognition of ATM as such a system was lagging and would lag for some time. AERA 3 would have been ahead of its time, which is a good lead for research but not an operational system.

To summarize, the idea for the TFM-ATC interface (and could have been extended to ASM) was almost the same as for any human-supervisory role over sector ATC in AERA 3: Each level decided about distinctly scaled objects with a statistical relation to subordinate objects that had their own focal-level decisions. The lowest levels have the most local and fast interactions and that is where automation belongs.

This idea could have exploited one of the more solid developments in AERA 3 by a colleague (Bill Niedringhaus) that was use of linear programming (LP) to define a set of planes (not airplanes) in the airspace that were the constraint set on permissible tracks [see William P. Niedringhaus, Automated planning for AERA 3 : maneuver option manager, Mitre/FAA, 1989, available online]. I wanted to

exploit the LP dual, the “shadow price” of the constraints on efficient tracks and trajectories. This became a concept of a price-like formal level to flight and track decisions. Like the “invisible hand” of the economy it would be an adaptive protocol form that would “guide” but not determine local decisions in the scale hierarchy. It could automate what TFM did through both sector track control and flight-traffic control. This idea never got farther than what I have stated.

## Reflections

This memoir has already stated my reflections on what happened to the NAS Plan, the AAS and AERA. At the time these reflections were intuitive for me and others but time has only sharpened the disappointment at what happened given the increasingly pressing issues of AVs and AI if not the still manual ATC and its interface with TFM.

The NAS Plan and AAS was widely indicted as too much of a "big bang". As I argued in its wake and now, that indictment is more the result of the changing views on "open" or "evolutionary" systems contemporary with the NAS Plan and AAS rather than inherent flaws in the Plan and the research path of AERA. It appears more obvious to me that the problems were the risk management of the program. The apparently rising costs and schedule slips were usual for such a program but the cost in particular was caught in a time of federal budget constraint. Despite our B-C analysis, evident to the CBO, Congress made no consistent comparisons of costs and benefits among programs. There was too much preconception of an investment in automation to reduce budget costs rather than to improve the ATC (or wider ATM) function of the airspace system. NextGen appropriated the larger mission, spoke the "architecture" language but was not the deliberate investment and research program that the NAS Plan was.

The focus on the AAS and AERA is not to ignore the many system sustainments and modernizations that were in the NAS Plan and implemented even in its lifetime to 1989. However, the conversion to the CIP emphasized more the hardware and less the functional improvements of AERA that the AAS was the platform for. The en route (ERAM) and terminal platform improvements (STARS) were a fragmentation and restart of the AAS, the ACF concept reviving only in 2011. STARS in particular had superior flight tracking RDP and last I heard was the candidate for incorporation into any consolidated platform. FDP remained an en route specialty and as I mentioned my late efforts to integrate that with terminal ATC failed.

I entitled a draft article on this subject "What was Wrong with the NAS Plan?" The answer I have given is its management and not its concept or scope. NextGen that was articulated ca. 2000 from the Free Flight concepts only accomplished drawing attention away from the "ground" or FAA part of the entire airspace and air transport system. The ATM system as Federal operation of a mode is unique and there may be a reversion to the highway-like concept of a public "infrastructure" for private commerce in which Federal or State regulation of use is contentious and ambiguous. This is being written (June 2024) as the congestion pricing plan for Manhattan is being pushed back. The capacity, financial, safety and generally ecological issues are still there. The unique concern for flight safety—that hardly reaches the crash costs on the highways—is hard to explain except for the "tragic event" bias. ITS however was a program in the other direction of bringing automation and communications into management of surface transport, but highways particularly. There is a path from SAGE to ITS via Mitre. Having worked also on ITS I can attest that there was no direct porting of the ATC-TFM concepts because cars are not aircraft. However more recent versions of the IVHS with AVs have already raised the issue of vehicle-to-other systems (V2X).

The most regrettable failure in the wake of the AAS and AERA is the discontinuity in research on the persistent issue of the human role in risk management, whether of projects or the operation of dangerous systems. There is an increase in complexity in the entire ecology because all interactions are extended and intensified when technology is applied to consumption of space or information. The ATC system has always been an avatar of that confrontation with complexity with high risk accountability that falls on humans. In the face of that complexity the tendency has been to recede from the kind of Federal research that was defense-motivated but with such remarkable impact on a cyber ecology. What is wrong with NextGen extends to all other systems with any public risk: It is not enough to envisage an "architecture" and then assume distributed-commercial developments will either implement all functional modules or adequately regulate the risks. We are still learning the responsibilities of an "infrastructure" that becomes ecological form. The increasing possibilities for humans with automation only confuse where the risk accountability lies.

\*\*\*\*\*

## **Attached Illustrations**

Picture Addendum to Gary G. Nelson: AERA 3 and the NAS Plan



Figure 1: NAS Stage A (ca. 1973), the mainframe generation in an Air Route Traffic Control Center (ARTCC). Picture from FAA at

[https://www.faa.gov/about/history/photo\\_album/air\\_traffic\\_control](https://www.faa.gov/about/history/photo_album/air_traffic_control)



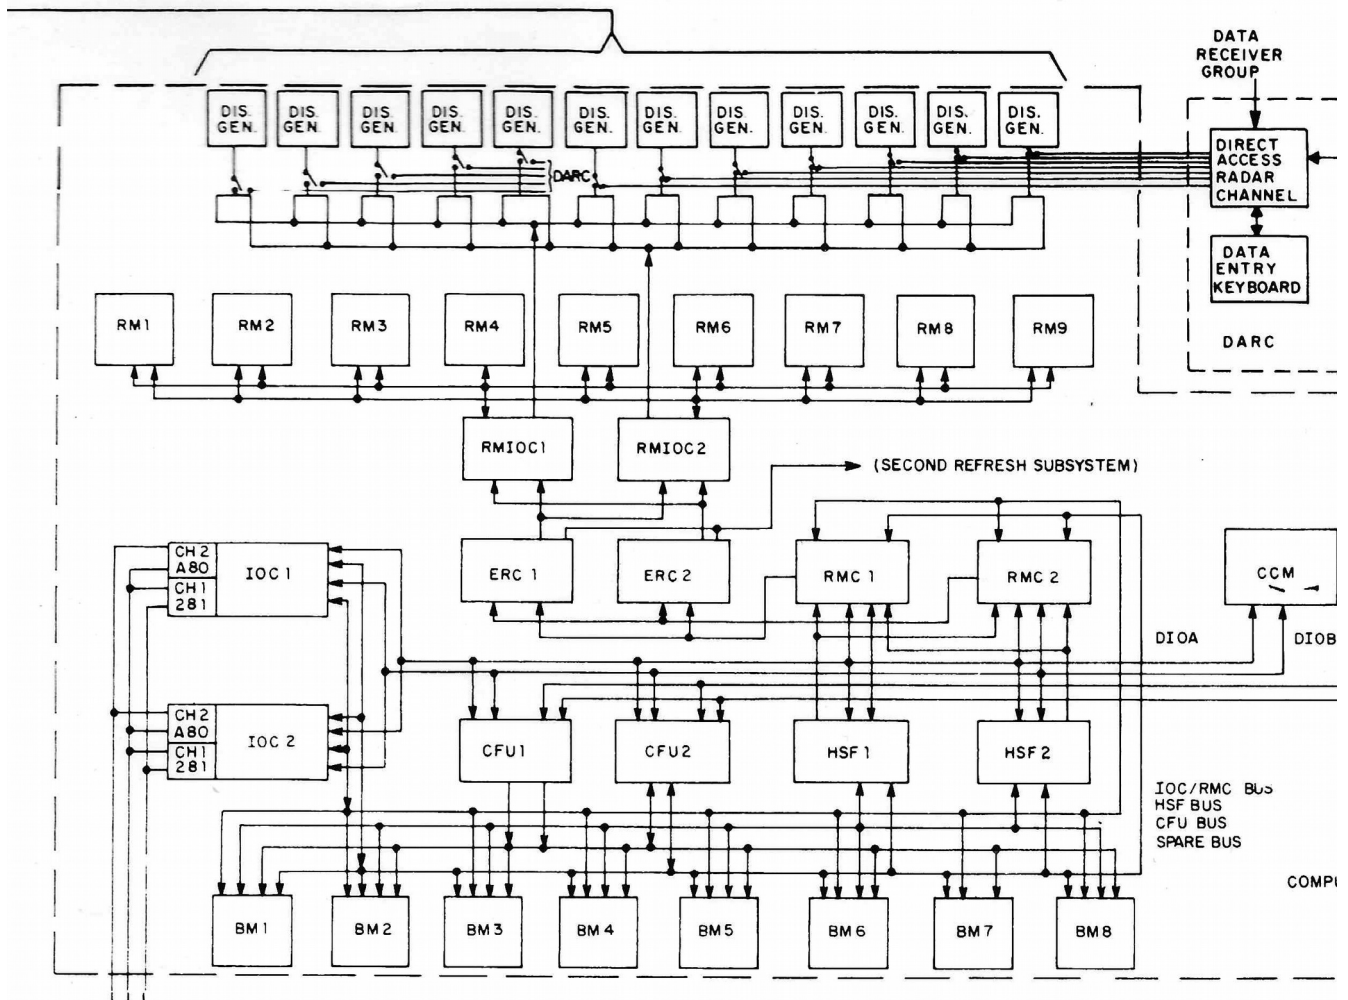


Figure 2: NAS Stage A (ca. 1973). Block diagram of the Computer Display Complex (CDC). This was the Raytheon "High Speed Switch" configuration. The alternative Display Computer Complex (DCC) used the IBM 9020 (System 360) processors similar in configuration to the CCC of Fig. 3.



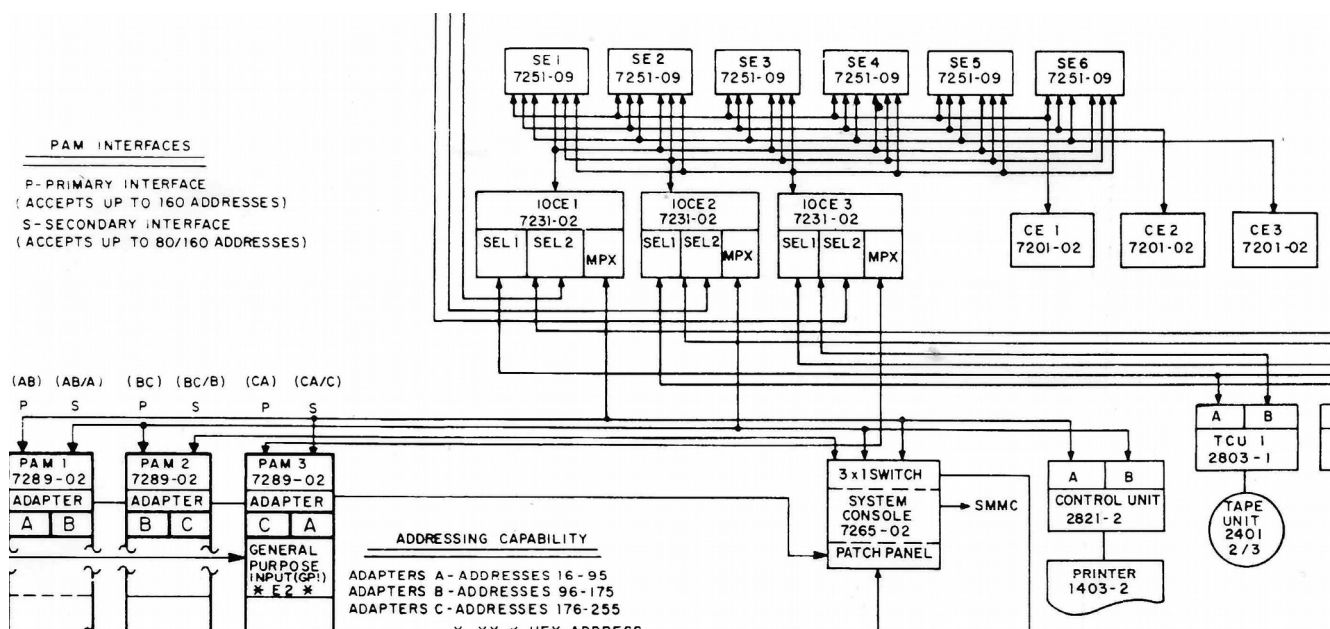


Figure 3: NAS Stage A (ca. 1973). Block diagram of the Central Computer Complex (CCC) also called the IBM 9020 system based on IBM 360 processors. Peripheral Adapter Module (PAM) handled external input-output to system excepting direct (wideband) surveillance. Post Host CCC (IBM 9672, after 1989) and Display Channel Replacement (DSR) (after 2000) configuration can be found at <https://www.faa.gov/documentLibrary/media/Order/6100.1H2.pdf>



Figure 4: NAS Stage A (ca. 1973). Pictures from FAA at [https://www.faa.gov/about/history/photo\\_album/air\\_traffic\\_control](https://www.faa.gov/about/history/photo_album/air_traffic_control). Above is Plan View Display (PVD) in En Route sector position. The original PVDs could be rotated horizontally to accommodate backup (digital channel failure) “shrimp boat” markers of flight identification. Below is the Computer Update Equipment (CUE) for sector participation in Flight Data Processing (FDP) but with paper flight strip interface visible in the strip bay behind the controller's head.