Defense Intelligence Reference Document

Acquisition Threat Support

12 January 2010

ICOD: 1 December 2009

DIA-08-0912-008

Materials for Advanced Aerospace Platforms

Materials for Advanced Aerospace Platforms

Prepared by:

(b)(3):10 USC 424

Defense Intelligence Agency

Author:

(b)(6)

Administrative Note

This product is one in a series of advanced technology reports produced in FY 2009 under the Defense Intelligence Agency, (b)(3):10 USC 424 Advanced Aerospace Weapon System Applications (AAWSA) Program. Comments or questions pertaining to this document should be addressed to (b)(3):10 USC 424;(b)(6), AAWSA Program Manager, Defense Intelligence Agency, ATTN: (b)(3):10 USC 424 Bldg 6000, Washington, DC 20340-5100.

Contents

Figures

Tables

Materials for Advanced Aerospace Platforms

Introduction

"Advanced aerospace platforms" is a broad topic that can be divided into several narrower subtopics to enable a more concise discussion of materials advances, challenges, and opportunities. Consequently, this document discusses the areas of launch vehicles, space vehicles, and space propulsion systems separately because their key requirements are often application specific, which affects materials selection decisions. In addition, single-use and reusable boosters have different durability requirements that directly impinge on design and materials selection. Furthermore, current engineering practice has evolved to the point that design synthesis must integrate the structure and construction materials to achieve optimum product performance. For example, the space shuttle was designed to meet customer-imposed mission requirements (range, payload, empty weight, landing capability, and so forth) without significant real-time consideration of materials capability. This approach led to significant compromises at later stages in the shuttle's development and maturation. (Arguably, the shuttle could be designed as a more efficient vehicle today.) In the extreme, a spectacular engineering failure was the National Aerospace Plane (also dubbed the Orient Express), which was launched as a military project and was intended to be a mach 12 reusable strike vehicle. This project rapidly became materials limited and was canceled in 1993, after about $750 million in federal R&D expenditures and a substantial private sector investment. The point is that any "clean sheet of paper design" must start with an assessment of the requirements for construction materials and be accompanied by a realistic assessment of the capability of currently available materials to meet these needs. If these two assessments indicate a gap between requirements and existing materials capability, a risk assessment and a risk-mitigation plan must be developed before expending engineering hours and funds.

Since the inception of manned space flight, the approach to design has changed to include the concept of damage tolerance. This shift in design philosophy was prompted by the (eventual) recognition that complex structures cannot be designed and produced with zero defects. With the maturation of fracture mechanics and means of reducing these concepts to practice, the transition from zero defects to defect tolerance became the norm. This new approach in turn led to recognition that high-performance materials required not only high specific strength and stiffness but pacing increases in strength with simultaneous improvements in fracture toughness and fatigue crack growth resistance. The introduction of damage tolerance was accompanied by a renewed emphasis on nondestructive inspection capabilities. This latter thrust was driven by the need to demonstrate the capability to reproducibly locate small flaws that could become failure initiation sites, either because of static or because of cyclic loading conditions. In the case of atmospheric flight, the U.S. Air Force has introduced standards for airframes (the Aircraft Structural Integrity Program) and propulsion systems (the Engine Structural Integrity Program) that tie structural life and reliability to

this demonstrated inspections capability. Implementation of these standards, starting with the B-1 bomber and the F-100 and F-110 engines, has dramatically reduced (but not eliminated) the incidence of catastrophic failures of critical components that endanger crews, vehicles, or both.

Taking these changes into account, this document outlines the current situation regarding the design and production of high-performance structures for aerospace platforms, including launch vehicles, space vehicles, and propulsion systems for transporting space vehicles (and payloads) into orbit.

Materials for Advanced Aerospace Platforms

LAUNCH VEHICLES

For the purposes of this document, launch vehicles are defined as the structure that supports and/or encloses the propulsion system, the fuel supply, and the crew or payload module. Launch vehicles today are either single use or multiple use after recovery and extensive refurbishment. This approach adds considerably to the cost of transporting a pound of payload into earth orbit, regardless of whether an unmanned satellite or a manned orbiting crew module that must withstand the temperatures and loads associated with safe reentry to earth. Furthermore, the larger the payloads are, the greater are the reaction forces the launch vehicle must withstand during launch. With the total weight of the payload, the empty weight of the launch vehicle, and fuel all needing to be lifted initially, fuel-efficient propulsion and lightweight launch vehicles are essential to maximizing the payload. Except in the area around the propulsion system exhaust, the temperatures experienced by launch vehicles during launch are not demanding. Therefore, advanced, high-strength aluminum (Al) alloys and polymer matrix carbon fiber composites (PMCs) are prime candidates for the parts of the structure that experience aerodynamic loads and where aerodynamic heating does not exceed about 125 °Celsius. One class of advanced Al alloys is the lithium (Li)-bearing alloys, such as Al alloy 2090. This alloy contains enough Li to reduce its density by 8 percent while increasing the elastic modulus (E) by 10 percent. Other, newer advanced Al alloys, such as 7050 and 2050, have been developed to have improved damage tolerance. These alloys have excellent specific strength at or near room temperature and experience no major loss of ductility at cryogenic temperatures. The newer variants of the 2000 and 7000 Al alloys also have substantially improved resistance to most types of corrosion, including exfoliation and stress corrosion cracking. This can be important in a reusable vehicle.

Perhaps the most important aspect of the improved Al alloys is their higher fracture toughness, accomplished through a combination of alloy composition control and improved processing. In alloy composition control, the concentrations of the residual elements iron (Fe), chromium (Cr), manganese (Mn), and silicon (Si) are reduced at the ingot stage. These elements combine with Al to form hard, brittle intermetallic compounds known as constituent phases. The advanced alloys contain fewer, smaller constituent phases, leading to improved fracture resistance and higher fracture toughness values. In applications such as body skins for commercial aircraft, this improved toughness has enabled an increase in the spacing of the circumferential fuselage frames, or "hat sections," that serve both as stiffeners and as crack stoppers to prevent a catastrophic failure during pressurization. For any given operating stress—in this case the pressurization stress—the spacing of the frames is directly related to the critical crack size of the body skin. Higher toughness alloys have larger critical crack sizes, and the frames can be spaced further apart without increasing the risk of catastrophic failure. The increased spacing ultimately allows a fuselage design that requires fewer frames. Consequently, the airplane benefits from a commensurate reduction both in weight and in manufacturing cost. Similar possibilities exist for the design of a fail-safe launch vehicle that has a lower empty weight. Clearly, the advanced Al alloys offer intrinsic improvements over the alloys used in the Saturn launch vehicle and introduce the prospect of new, more efficient launch vehicle designs.

The newer Al alloys also can be specially processed to render them superplastically formable. This capability opens a realm of possibilities to replace structures that, in the absence of this capability, are machined from thick plate. Very large structures are produced in sections that must be joined. Conventional fusion welding techniques do not work for high-strength Al alloys such as 7075, 7050, 2024, or 2050 because either the welds lead to cracks or the welds made under conditions that avoid cracking have greatly reduced tensile properties. Because these alloys are not amenable to welding, heavier, fatigue-prone mechanically fastened joints must be used. Recently, scientists developed a joining process that permits joining of Al alloys such as 7050. This process, called friction stir welding (FSW), allows joint designs in a variety of configurations that were not considered possible when fusion welding was the only alternative.

In FSW, a rotating steel tool is inserted into the seam between the two Al alloy pieces to be joined. As the rotating tool is driven forward, the friction between the tool and the work piece generates enough heat to soften the Al alloy without melting it. A schematic of this process is shown in Figure 1.

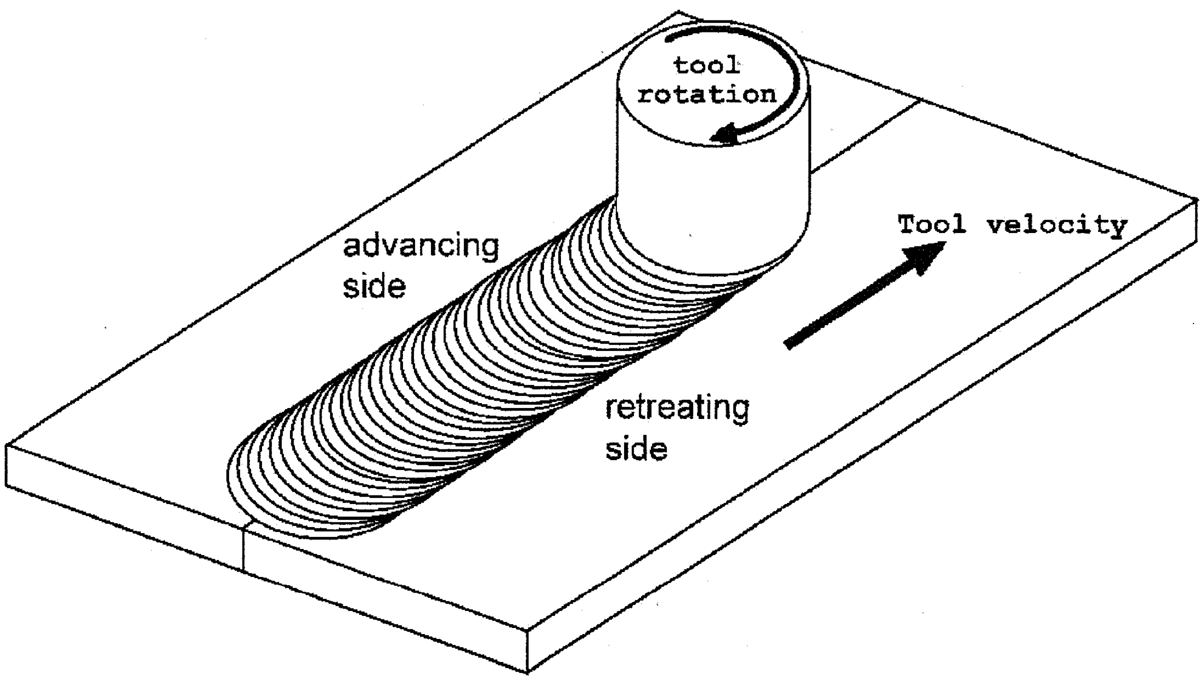

Figure 1. Schematic Diagram of Friction Stir Welding

The combined action of the rotation and the traversing of the tool essentially kneads the two pieces together, leaving a mechanically sound joint. Although the weld properties may be somewhat inferior to those of the base metal, they are good enough that a relatively small increase in thickness at the joint position can compensate. Although substantial development of the FSW process is ongoing, FSW already has been put into practice. For example, the current external propellant tank on the space shuttle is made from an Al-Li alloy fabricated through FSW. The weight advantage of using welded as opposed to bolted joints in a large structure such as a launch vehicle is considerable. With earlier high-strength alloys such as 7075, concerns about fracture toughness in conjunction with monolithic structures would have caused a welded construction to be considered too risky. Today, the combination of higher toughness alloys and FSW opens up the possibility of greater design flexibility resulting in lighter large structures with equal or greater reliability than earlier ones.

In sum, metallic, nonreusable (at least nominally so) launch vehicles made from advanced Al alloys and fabricated through FSW constitute an incremental but significant improvement over earlier versions.

In recent years, PMCs have matured significantly. For many components that are not exposed to elevated temperatures, PMCs provide a degree of design flexibility not readily available in metals. Consequently, PMC materials have begun to supplant Al alloys in the construction of commercial subsonic aircraft. The use of PMCs in the empennage of the Boeing 777 was one of the first examples of Al alloys being displaced. Subsequently, the new Boeing 787 has more structure made from composites than from metallic materials. Once PMCs are introduced into a structure in significant quantities, a constraint related to galvanic incompatibility between the PMC structure and any adjoining Al alloys also is introduced. When a PMC structure is in direct contact with an Al alloy structure, catastrophic corrosion of the Al alloy components can occur. In the Boeing 787, the remedy for this concern is the use of titanium (Ti) alloys in areas where there is direct contact between the metallic and the PMC structures. This is directly analogous to the plastic bushing a plumber puts in the joint between copper and iron piping. Notwithstanding this constraint, the specific strength and stiffness of PMC structures make a compelling argument for their application in high-performance structures, such as launch vehicles.

Composite structures can be manufactured using one of three methods: hand layup of pre-preg, automated tow placement, and resin transfer molding.

- The most rudimentary of these, but also the most flexible, is hand layup of pre-preg. This method uses sheets of material that contain both the fiber and the polymeric matrix (called pre-preg). The polymeric matrix can be either a thermoset (for example, epoxy) or a thermoplastic. Individual plies are cut from the pre-preg typically using a numerically controlled laser or mechanical cutting device and are laid up to form the desired shape. Areas that have heavier loads contain more plies locally, and the plies are cut in an orientation with respect to the fiber direction in the pre-preg to achieve the desired strength relative to the principal load path. These plies are carefully placed according to a drawing (blueprint), making hand layup a labor-intensive process and, therefore, making parts made using this method expensive. During ply placement, it is critical that no ply wrinkles are introduced, as these create severe reductions in the local load-bearing capability of the final component. Once all the plies are in their proper places, the article is placed in a vacuum-tight bag that is evacuated and placed in an autoclave for curing of the epoxy matrix or fusing of the thermoplastic. A disadvantage of a pre-preg whose matrix is a thermoset is limited shelf life. In practice, this is managed to a degree by storing the pre-preg in a freezer to slow the rate of chemical reaction that sets the epoxy. However, this does not completely halt the reaction, causing these materials to have a shelf life beyond which they are not easily manipulated during layup and do not develop full strength after curing in the autoclave. An additional issue is out time—the time the pre-preg can be out of the freezer during layup before the reaction proceeds at an accelerated rate and reaches a point at which the material is not suitable for the reasons stated earlier. Clearly, the time required for layup places practical limitations on component size.

- In automated tow placement, thin ribbons of a pre-preg are fed off a drum or rolled into a computer numerically controlled machine that places them in the desired position. In principle, this process trades recurring labor cost for up-front capital investment (the tow placement machine) and programming time. If the anticipated volume of identical parts is high enough to amortize the capital investment and, particularly, the programming cost, this can be an attractive means of reducing manufacturing costs. For axisymmetric shapes, such as cylinders, this essentially becomes a winding process and is quite efficient. An example of a finished composite fuselage barrel section for the Boeing 787 is shown in Figure 2. For more irregular three-dimensional shapes, such as a spar or a strut, placing the tows becomes much more difficult and presents a fundamental limitation. Consequently, PMC structures with complex shapes are still for the most part made using the hand layup process. A variant of automated tow placement is compression, whereby a preform, made by automated tow placement, is forced by a press into a preshaped die. This process allows fabrication of more complex shapes, but the rigidity of the fiber and the extreme anisotropy of the tows can lead to wrinkles, which are not acceptable because of the reductions in properties these cause.

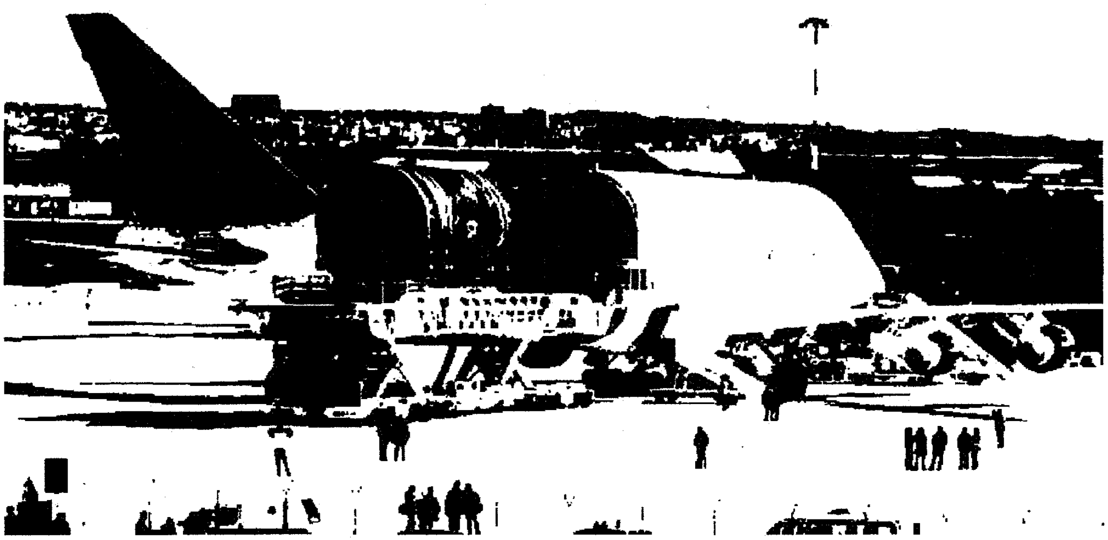

Figure 2. Specially Modified 747 Transporter Unloading a Boeing 787 Composite Fuselage Barrel Section

- The third main composite fabrication method, resin transfer molding (RTM), begins with a woven fiber mat or preform. The polymeric matrix is injected into this mat to create a fully dense composite. The major benefit of RTM is that it permits use of a three-dimensional weave that minimizes the risk of delamination between plies. Note that, with the other two methods, the material is reinforced in only two dimensions (the plane of the pre-preg or tows). RTM's limitations include the viscosity of the resins used. If the resin is too viscous, injecting it will either distort the fiber architecture of the woven preform or not fully penetrate the preform, leaving voids. This limitation makes it difficult to use many thermoplastic resins that are otherwise attractive because they are recyclable and have much longer shelf lives than thermosets (epoxies). Another, more obvious limitation of RTM is component size. This is in part because weaving of very large fiber preforms is challenging and requires a very large weaving machine. It also requires a large injection machine capable of multiple injection sites to ensure the complete infiltration necessary to avoid formation of voids.

In practice, there are several significant challenges associated with manufacturing large PMC structures today. Among these, perhaps the most significant are manufacturing cost and the difficulty associated with making large, nonaxisymmetric components. Large axisymmetric components that can be produced by winding tows of pre-preg are clearly possible today, as demonstrated by the Boeing 787 fuselage whose section is shown in Figure 2. For other shapes, because the load-bearing capability of PMCs depends on the transfer of externally imposed loads to the strong, stiff fibers, joints that intersect the principal load path become problematic. This is because the fibers and, as a direct consequence, the load path are discontinuous, requiring local section size adjustments to offset this local structural inefficiency. Thus, achieving the most structurally efficient use of PMCs requires monolithic structures with continuous fibers. In large structures, this is at best a challenge. Even in the Boeing 787's PMC fuselage, mechanically fastened joints are used to connect adjacent barrel sections. In the fuselage, as in any cylindrical pressure vessel, the principal stresses are hoop stresses stemming from pressurization during flight. However, the compromise in structural efficiency is minimized by the circumferential orientation of the joints. Although these bolted joints add weight, the overall structural efficiency of the PMC structure is still better, albeit considerably more expensive, than a longitudinally and circumferentially stiffened Al structure. The circumferential joints also create discontinuous longitudinal crack paths that improve the structure's damage tolerance. Some of the added expense of the PMC fuselage stems from the use of Ti alloy fasteners because of the galvanic coupling issues that would accompany Al or steel fasteners.

For other applications, which are limited by different material properties, the PMC system can be tailored to optimize structural performance. This is possible because both the matrix and the fiber can be independently selected. Moreover, the fiber "architecture" (fiber orientation, weave geometry, and fiber volume fraction) can be varied spatially to optimize load-bearing capability under complex stress states. For example, again drawing on recent applications in commercial aircraft, the fan blades of the large, high-bypass-ratio turbofan engine (GE90) produced by General Electric for the Boeing 777 are made from PMCs. The limiting design consideration for these fan blades is resistance to bird strikes. To optimize the PMCs' impact resistance, a medium-modulus, high-tensile-strength carbon fiber was selected in combination with a thermoplastic toughened epoxy matrix. Furthermore, the fiber architecture was set to optimize the bending strength under the impact of a bird. The GE90 fan blades are produced by hand layup and are quite costly to produce. In the 10-plus years that these fan blades have been in service, not a single unscheduled engine removal related to the PMC fan blades has occurred. The competitive fan blade technology for B-777–class engines is hollow Ti, which is used by both Pratt & Whitney and Rolls Royce. By all informal accounts, these blades are cost intensive. This example supports the unwritten rule that the pathways leading to high-performance, high-value structures typically are not technology dependent but are cost intensive, no matter what technology is employed to meet the requirements.

An additional characteristic of PMC structures—one related to the laminated construction of components made by hand layup or automated tow placement methods—is their susceptibility to formation of delaminations when impacted perpendicular to the plane of the plies. This is due to the mismatch in bending stiffness between adjacent plies that have different unidirectional fiber orientations. This mismatch causes shear stresses to develop that can exceed the shear strength of the interlaminate bonds, causing small, embedded cracks to form. Under subsequent inplane compression loading, the laminates bow because of the Poisson stresses and separate because of the lack of an interlaminar bond to hold them together. In significant compression loads, the laminates buckle, and this leads to structural failure. This phenomenon, called compression after impact, is an insidious failure mode because the delaminations are not externally detectable unless ultrasonic inspection methods are used. Sources of such an impact include dropped tools, foreign objects (for example, meteorites), and, perhaps most commonly, hail storms. Ultrasonic inspection, if required, is expensive and time consuming. The latter concern in turn affects vehicle availability and turnaround time. PMC parts made using RTM typically have reinforcing fibers in the through-thickness direction, so concerns about compression after impact are minimal.

The maximum temperature at which PMCs can be used is limited by PMCs' susceptibility to oxygen degradation of the polymeric matrix. The maximum-use temperature for prolonged exposure is determined by the thermal oxidative stability (TOS) of a particular resin. The TOS, like any chemical reaction, is determined by both time and temperature. The glass transition temperature (T_g) of the polymeric matrix also imposes strength and dimensional stability limitations independent of the TOS limits. This is particularly true for thermoplastics. However, the TOS limits usually impose lower temperature limits than T_g if prolonged thermal exposure is contemplated. Three distinct classes or groups of resins exist, each with a successively higher temperature capability. These are conventional epoxies and most thermal plastics, bismaleimides (BMIs), and linear polyimides. The first class is limited by TOS to about 125° Celsius. BMIs can be used to about 175° Celsius. Linear polyimides, such as the in situ polymerization of monomer reactants (PMR) group of thermosetting formulations, can be used to about 300° Celsius. Many of the PMR resins contain the hazardous compound *methylenedianiline*, which requires special care during use, including protective clothing to limit personnel exposure (for example, during ply cutting and hand layup). This requirement reduces productivity, adds cost, and creates a degree of liability concern for the manufacturer of the PMC components. For RTM, the BMIs and PMR resins typically have higher viscosity and require commensurately higher injection temperatures to reduce the viscosity to manageable levels for reasons discussed earlier. New resins are constantly being developed, including some with attractive properties. However, many of the sources of these resins are startup companies that exist on R&D funding, often in the form of U.S. government SBIR (Small Business Innovative Research) projects. Such companies are good at innovation but often have limited experience transitioning new products from the laboratory to large-scale production. Furthermore, in the current economic climate, access to sufficient capital to set up production-scale capacity can be a formidable problem for a small company. Larger companies (for example, BASF, DuPont, GE Plastics) typically are not interested in materials with small annual sales volume, so even licensing the new material to one of these companies may not be commercially feasible.

This foregoing discussion highlights a commonly encountered inconsistency between technical innovation and commercial progress. Until a production-scale source exists, pricing of new materials is at best highly uncertain and potentially unstable over time. This is in part due to the uncertainty surrounding demand and the associated volume of material that will be required. Taken together, these factors act as a clear deterrent to the adoption of new materials of literally all classes (polymers, metals, and ceramics).

For heavily loaded structures or structures that will experience temperatures higher than about 200 °Celsius, Ti alloys are the preferred material class. Ti alloys are about half as dense as steel or Ni-base alloys and possess a density-corrected strength and stiffness competitive with that of other metallic materials. Numerous grades of Ti alloys are in use today, and a strong domestic industrial base of suppliers exists for nearly all these grades. Ti alloys for structural applications can be divided into three groups based on their metallurgy: near α alloys, $\alpha+\beta$ alloys, and metastable β alloys (commonly called β alloys for short). Both the aircraft and propulsion original equipment manufacturers and several private engineering firms have extensive design experience with Ti alloys. Although most of this experience is with Ti-6Al-4V (Ti-6-4), other alloys also are widely used, particularly in jet engines and liquid-fueled rocket engines. The alloy most likely to be used in a launch vehicle, Ti-6-4, has been available and in use for more than 40 years but is still highly competitive with newer grades in large part because of its versatility. Potential applications of other alloys that have particularly attractive characteristics, such as a higher temperature capability, are discussed in the Reusable Reentry Vehicle and Propulsion Systems sections of this document.

The most commonly used structural grade of Ti alloy today, Ti-6-4, can be readily fusion welded, formed both by forging and as a sheet product, and conventionally machined, although each of these operations requires special precautions. In addition, the feasibility of friction stir welding has been demonstrated. Ti-6-4 also can be superplastically formed and diffusion bonded, enabling synthesis of innovatively shaped components. For launch vehicles, the most likely application for Ti alloys is in the structure that carries the reactions from the propulsion system to the vehicle itself. These applications typically involve heavy sections to accommodate the large loads, and the main property requirements are high strength, fatigue resistance, and fracture toughness. Ti alloys have been used in both military and commercial aircraft. Heavily loaded components in service today include the wing box of the B-1B bomber (Ti-6-4), the landing gear beam in the B-747 (Ti-6-4), and the landing gear truck beam in the B-777 (higher strength Ti-10V-2Fe-3Al [Ti-10-2-3]). The choice of Ti-10-2-3 for the landing gear truck beam reflects the time-based maturation of the newer β-Ti alloys, such as Ti-10-2-3. These alloys have the advantage of being "deep hardenable"—compared with Ti-6-4, they can develop full strength in thicker sections during heat treatment. For example, that the B-777's truck beam is up to 6 inches thick in some locations factored significantly in the choice of Ti-10-2-3. Because Ti alloys are about half as dense as steel, they are very competitive on a density-corrected basis. However, mass is not the only driver for some applications; the volume of a component also must be compatible with the space available for it. This factor also is a consideration in the choice of Ti-10-2-3 for the B-777 truck beam. And it also applies to an application such as landing gear, because it is retracted into the fuselage for aerodynamic reasons.

β-Ti alloys would be an attractive option for a launch vehicle with a very heavy payload. These alloys can be processed and heat-treated to ultimate tensile strengths greater than 1,300 mega-Pascals (MPa), making them very efficient structural alloys. However, at these strength levels, the fracture toughness is decreased to about 40 MPa-$m^{0.5}$, making damage tolerance marginal. For example, using a design stress that is two-thirds that of ultimate tensile strength, the critical crack size for an alloy with these properties is about 4 millimeters. Such a small critical crack size poses a challenge to any required field inspections associated with reusability requirements. β-Ti alloys' stiffness also is as much as 10 percent lower than that of α+β alloys. Ten percent lower stiffness generally is not an issue in a tension-loaded structure, but it can be an issue for a compression-loaded structure because of the potential for buckling. Designs with a higher section modulus can eliminate this concern, but the additional shape complexity will almost certainly add cost. In extreme circumstances, Ti alloys can be reinforced with ceramic fibers (typically silicon carbide) to increase their intrinsic stiffness, as will be discussed later.

Ti alloys are reactive when exposed to air at temperatures of 550 °Celsius or above. Consequently, any manufacturing operations that exceed this temperature limit must be performed in a protective atmosphere of argon (Ar) or helium (He) gas. An exception occurs during forging if enough excess material is left on the raw forging to contain the oxygen-contaminated layer so it can be machined away during the final machining of the finished component. How to deal with this reactivity issue is well understood, and it poses no concern other than the additional costs associated with the excess material and additional machining. In other operations, such as welding, use of specially designed fixtures incorporating inert shielding gas also effectively eliminates concerns about oxygen contamination. In hot-forming applications, the as-formed part is typically chemically milled to remove the oxygen-rich surface layer because this layer typically has lower ductility and can cause fatigue cracking in service. The practice of eliminating all oxygen-contaminated material has served the aerospace industry well over the years but is quite conservative and restrictive. This matter is discussed in greater detail in the Reusable Reentry Vehicle section of this document.

In summary, although the choices of materials for launch vehicles are in principle numerous, in practice these choices are reduced by a variety of application-specific considerations that include manufacturing capability for large components and manufacturing cost. The foregoing discussion has attempted to examine the prospects for advanced Al alloys, PMCs, and Ti alloys in light of these perceived practical constraints. On a case-by-case basis, a variety of requirements are imposed by design constraints, which some materials meet more readily than others. In all cases, discussion of available material options was constrained by the assumption that the maximum service temperatures would be relatively low. Consequently, the material classes discussed here are all intended for relatively low-temperature use. The separate case of a reusable single-stage-to-orbit vehicle, where operating temperature requirements are considerably higher, is discussed in a later section of this document. Ultimately, materials are selected to optimize structural performance, and a coordinated approach of materials selection and geometric design is essential to this. Going forward, a design using a synthesis process that treats form, fit, function, and materials capability as equal constraints is needed to achieve true optimum structural efficiency.

REUSABLE CREW MODULES

The concept of manned orbital crew modules has evolved from the Mercury capsules to the Gemini and Apollo programs to the space shuttle, the first fully reusable crew module. The shuttle also has a combined payload bay used for transporting satellites into orbit and hardware for developing the International Space Station and for repairing and refurbishing the Hubble telescope, among other uses. The reusable nature of the shuttle crew module introduced a number of design and materials selection challenges. Perhaps foremost among these is the requirement for a thermal protection system (TPS) that would protect the crew during reentry and also minimize the intermission refurbishment requirements of the spacecraft itself. During the early days of the shuttle development program, there was much initial interest in a metallic TPS because it appeared to better meet the program's needs. Ultimately, however, ceramic tiles were used on the underbody and carbon-carbon composites (C-CCs) were used on the leading edges of the wings. The shuttle design itself can therefore be characterized as a "cold structure" with an insulating TPS. For example, much of the shuttle load-bearing structure is made of the Al alloy 2219, in part because it is fusion weldable and in part because it retains its strength at moderately elevated temperatures better than other high-strength Al alloys can. The refurbishment needs of the ceramic shuttle tiles after each flight reputedly are considerable and increase with vehicle age. The C-CC wing leading edges are basically not repairable but require scrutiny. Hindsight shows that C-CCs "age" and lose much of their fracture toughness during repeated thermal exposure. (This loss of toughness was a prime factor in the Columbia disaster. Had a metallic heat shield that included the wing leading edges been used, this disaster arguably could have been avoided.) If the discussion of a metallic TPS were held today, the outcome might not be much different. Certainly any serious consideration of a reusable single-stage-to-orbit vehicle today would need to reopen the discussion of a metallic TPS. The challenges and opportunities associated with a metallic TPS are discussed later.

The design efficiency of an integrated TPS and load-bearing structure is extremely attractive. Such a design requires availability of high-temperature alloys that also have good strength at the moderate temperatures to withstand the aerodynamic and vibrational loads encountered during launch and orbital insertion. Any attractive alloy also must have reasonable intrinsic resistance to oxidation at the reentry temperatures and should be capable of being fabricated into sheet gauges at reasonable cost. Meeting these various requirements in combination becomes quite daunting. Earlier programs such as the DynaSoar reusable reentry glider devoted considerable time and resources to examining the use of refractory alloys such as Mo-0.5%Ti (Moly half Ti) and several Columbia-based alloys for the TPS and some hot structure. All the refractory metal alloys are solid solution strengthened and, consequently, have relatively low ambient temperature strengths. They also are quite dense, making the density-corrected strength even less attractive. Refractory metals and their alloys react extensively when exposed to air at elevated temperatures. Therefore, even if the mechanical property limitations could be overcome, any hot structure would require an oxidation-resistant protective coating. In the case of the DynaSoar program, scientists extensively investigated a surface conversion coating of $MoSi_2$ formed by reacting the Moly half-Ti alloy with Si powder in a high-temperature fluidized bed. With the benefit of hindsight, it is now unclear how a large structure could have been successfully coated in this manner. This is particularly true when the brittle nature of $MoSi_2$ is considered in light of the thermal stresses that were certain to develop in a large structure placed in the fluidized bed.

Reusable Single-Stage-to-Orbit Vehicles

A reusable single-stage-to-orbit (SSO) vehicle will require metallic materials for the TPS and for much of the other hot structure. This will be challenging from the standpoint of an empty vehicle weight. With the exception of military applications, which are outside the scope of this document, empty vehicle weight is a critical metric because every additional pound of empty weight reduces the payload by the same amount (assuming a fixed propulsion capability). As mentioned earlier, the key to a lightweight vehicle is the use of design methods that integrate the TPS and the load-bearing structure to minimize structural redundancy and single-function structure (for example, a TPS that is not load bearing). Achieving this goal will require new design paradigms that incorporate true synthesis of new structural concepts. In reality, such designs can be completed only if they are based on a detailed set of mission requirements, including the number of missions and expectations for turnaround time between missions. Furthermore, operational parameters such as the value of inserting a pound of payload into orbit are needed to bound the cost of the initial vehicle and the maintenance cost per mission (translated into cost per pound of payload). Absent such specific data, the following discusses possibilities for materials systems that can enable a reusable SSO vehicle. It is perhaps more efficient to discuss these materials according to their principal capability and the anticipated temperature regime in which they can be used most productively. This categorization method is illustrated in Table 1.

Table 1. Potential Materials by Use Temperature Regime and Property

Temperature Regime	Specific Strength	Specific Stiffness	Fatigue Resistance	Fracture Toughness	Creep Resistance
Ambient up to 250 °C	PMCs; Advanced Al alloys; AMCs*; Ti alloys; TMCs*	PMCs; AMCs; TMCs	PMCs; AMCs; Ti Alloys; TMCs	PMCs; Ti alloys	Ti alloys; TMCs
250 °C up to 550 °C	Ti alloys; TMCs	TMCs	Ti alloys; TMCs	Ti alloys	Ti alloys; TMCs
Above 550 °C	Ni-base alloys; Ti aluminides; Refractory metal alloys; CMCs*; C-C Cs*	Ti aluminides; CMCs; C-CCs	Ni-base alloys; Refractory metal alloys	Ni-base alloys; CMCs; Refractory metal alloys	Ni-base alloys; Ti aluminides; CMCs; Refractory metal alloys; C-C Cs

AMCs = Al matrix composites; TMCs = titanium matrix composites; CMCs = ceramic matrix composites; C-CCs = carbon-carbon composites

Table 1 shows there are essentially 10 distinct classes of candidate materials for use in a reusable SSO vehicle. Their usefulness for specific applications and components depends on the operating temperature regime and the design-limiting material property. Some background and characteristics of each of these materials are described in this section. Applications of some of these materials have already been addressed, so the discussion here is limited to applications for a reusable SSO vehicle.

Advanced Al Alloys

This class of materials is limited by its temperature capability because, even if used in an embedded structure, the thermal soakback from the hot external structure can lead to softening in real time or overaging during extended exposure (that is, after several missions). Scientists have worked to develop high-temperature powder-metallurgy Al alloys, mainly alloys containing Fe, cobalt (Co) and cesium (Ce) or Mn and Si. These alloys are not routinely produced today owing to a lack of demand stemming in part from their cost and their relative lack of maturity as judged by the high variability in mechanical properties between lots of material. Consequently, the main application for Al alloys is in the crew compartment itself, where temperatures must be maintained at levels that are tolerable for the human occupants. Al alloys have marginal stiffness, even on a density-corrected basis. However, using fabricated panels with Al face sheets and a lightweight core can minimize this limitation. Such panels have a high section modulus, which increases the structural stiffness without adding much weight. Earlier uses of phenolic honeycomb cores experienced only limited success because of the tendency of the core to absorb water from the environment if a breach of the face sheet or the face sheet core bond occurred. Today, Al phenolic honeycomb sandwich construction is unpopular among aircraft and spacecraft designers because of this earlier, unfavorable experience. However, a new possibility for lightweight stiff structures—Al face sheets with a foamed Al alloy core—is worthy of consideration. Considerable progress in making uniform-density Al and other metallic foams has been realized in the past 10 years. This new class of sandwich materials is sufficiently different from the earlier versions to merit a careful evaluation.

Polymer Matrix Composites

As described in the Launch Vehicle section, polymer matrix composites have excellent strength, stiffness, fatigue resistance, and fracture toughness. All polymeric materials are to some degree susceptible to degradation when exposed to ultraviolet (UV) radiation. For an orbiting vehicle, a major limitation is the effects of UV radiation on the polymer matrix. Given their limited time at altitude, this vulnerability is not an issue for launch vehicles; however, it is a concern for structures that remain in orbit for an extended time at altitudes where the UV intensity is much greater. Although there are coatings that protect substrates (PMCs in this case) from UV radiation, the risk of these coatings being breached and the uncertainty about their effectiveness in orbit make UV exposure an ongoing concern. Any consideration of using PMCs for exterior applications would need to include an evaluation of their effectiveness. As discussed for Al alloys, PMCs also could be considered for interior applications where temperatures are within the material capabilities. As discussed earlier, some polyimide resins are usable up to about 300 °Celsius; therefore, the low density and excellent specific stiffness of PMCs make them viable candidates for some components. As also discussed, some polyimides, such as PMR-15, contain *methylenedianiline*, and care must be taken to ensure that any outgassing that occurs in space would not liberate any of this toxic material—at least not where it could be ingested by the crew.

Al Matrix Composites

Al matrix composites (AMCs) are typically made by mixing short fibers or even particulate silicon carbide (SiC) with Al alloy powder and hot-pressing or mechanically consolidating the mixture by extrusion or forging. AMCs have higher modulus and strength but suffer ductility losses owing to the hard, nondeforming SiC second phase. They also are much better in fatigue because the SiC particles or short fibers mitigate the effects of planar slip in precipitation-hardening alloy matrices. Machining, fusion-welding, or cold-forming AMCs is difficult. Mechanical fastening is possible, but drilling fastener holes is difficult and expensive because the wear caused by the hard, reinforcing phases shortens tool life.

Most of the AMCs produced to date have used relatively simple Al matrix alloys such as 6061 and 5083. This presumably is because of the conventional wisdom that AMCs are not heat treatable and the perception that no benefit is gained by using more complex alloy matrices. This is probably the case for a conventional solution treat-and-age approach. Consequently, there has been little effort to optimize AMC systems. For the right application where AMCs could provide a significant benefit, this could present an opportunity. For example, using the high-temperature Al-Fe-Co-Ce alloy powder as the matrix could prove interesting and might permit use of AMCs at up to 200 °Celsius—a temperature at which Ti alloys do not provide any significant advantage, but one that is too high for conventional Al alloys to be suitable.

In sum, while AMCs are not really a commercial materials system today, sufficient research has been performed to establish proof of concept. If an adequate market for AMCs were to emerge, the time and cost to make them commercially available could prove acceptable.

Ti Alloys

Ti alloys also have been discussed earlier, but mainly in the context of heavy-section, large-load-bearing applications for launch vehicles. Here, the potential of Ti alloys for lighter gauge applications in the warm structure and the TPS is considered. Table 1 shows that at intermediate temperatures, Ti alloys are attractive in all aspects except for specific stiffness. What this table does not capture is Ti alloys' propensity to react with oxygen in the air to form an oxygen-stabilized α phase layer on the surface known as α case. An example of α case is shown in Figure 3.

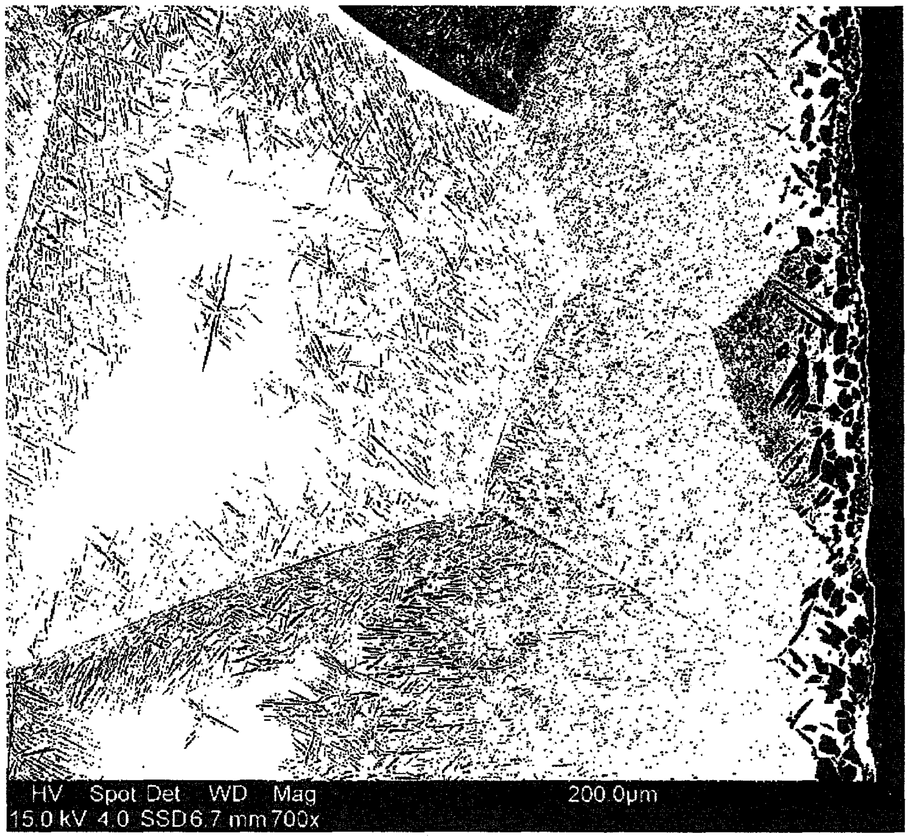

Figure 3. Micrograph Showing α Phase Formation at the Surface of a Ti Alloy That has Been Exposed to Air at Elevated Temperature

The α case is harder than the matrix because oxygen is a potent α-phase solid-solution strengthener. As with almost all other strengthening reactions, the increased strength is accompanied by reduced ductility. Again, drawing on conventional wisdom, the presence of α case in sheet structures has been forbidden by specification, design practice, or whatever means a company uses to manage its hardware. Less clear is how truly detrimental α case is to properties. Essentially no effort has been made to determine whether it can be tolerated if the affected hardware is allowed to operate at a modestly reduced stress. The industry standard for α case has essentially been one of zero tolerance. Given the significant potential weight advantage associated with use of Ti alloys in portions of the TPS and warm structure, this conservative approach needs to be revisited. Several key questions related to this are:

- Is α case truly detrimental to the load-bearing capability of Ti alloy sheet structures?
- If so, is there a limiting amount that can be tolerated without significantly degrading the structural capability?
- Which properties are the most severely degraded?
- Are α case formation and property degradation alloy dependent?
- If so, which alloys are the most tolerant of α case formation?

- Once α case is present, can the structure be repaired (for example, by fusion or friction stir welding)?

The Air Force Materials and Manufacturing Directorate is starting a new project intended to address these questions. The motivation for this program is hypersonic flight vehicles. The results from this U.S. Air Force program should prove highly useful to the design of future reusable SSO vehicles.

Earlier uses of Ti alloys at high temperatures included the skin and much of the load-bearing structure of the SR-71 Blackbird. This airplane flew successfully at peak speeds in excess of mach 3.2 for 34 years (1964-1998). While the maximum skin temperatures are not readily available, they were in excess of 300 °Celsius. There were no known issues involving α case during the SR-71's service. Notably, the primary alloy used for the SR-71 was one of the original β–Ti alloys, B-120 VCA, the composition of which is Ti-13V-11Cr-3Al. The primary reason for choosing this alloy was that it is much easier than any of the α+β Ti alloys are to roll into sheet gauges. Today there are newer β–Ti sheet alloys featuring a better balance of properties that could be used in the same way as B-120 VCA. The most common of these is Ti-15V-3Cr-3Sn-3Al. However, the successful use of B-120 VCA raises the question of whether β–Ti alloys are more resistant than α+β alloys such as Ti-6-4 are to α case formation.

The attraction to using Ti alloys, in addition to their structural efficiency, is the extensive industrial base for making the material in a variety of product forms and the extensive knowledge base resulting from the many successful applications of Ti alloys in high-performance products. For example, the ability to superplastically form Ti alloys such as Ti-6-4 creates the opportunity for design of a structure that functions both as load bearing and as thermal protection.

Ti Matrix Composites

As Table 1 showed, Ti alloys are not especially attractive for their specific stiffness. One way to overcome this limitation is to reinforce a Ti alloy matrix with SiC fibers. In this case, the fibers are "long" fibers—they have sufficient length for the matrix to transfer the maximum possible fraction of the external load to the fiber. The fibers are essentially monofilaments and must be carefully placed so adjacent fibers do not touch one another. Areas of contact between fibers essentially are incipient cracks that degrade the mechanical strength. As Table 2 shows, TMCs have excellent properties.

Table 2. Example of Properties of Ti Matrix Composites

Property	Property Value (English /Metric Units)
Ultimate tensile strength	276 ksi / 1902 MPa
Young's modulus	32.8 msi / 226 GPa
Strain to fracture	0.95%
Density	0.16 lb/in^3 / 4.43 g/cm^3
Fiber Volume fraction	0.39

Figure 4 shows an example of a TI matrix composite (TMC) cross section.

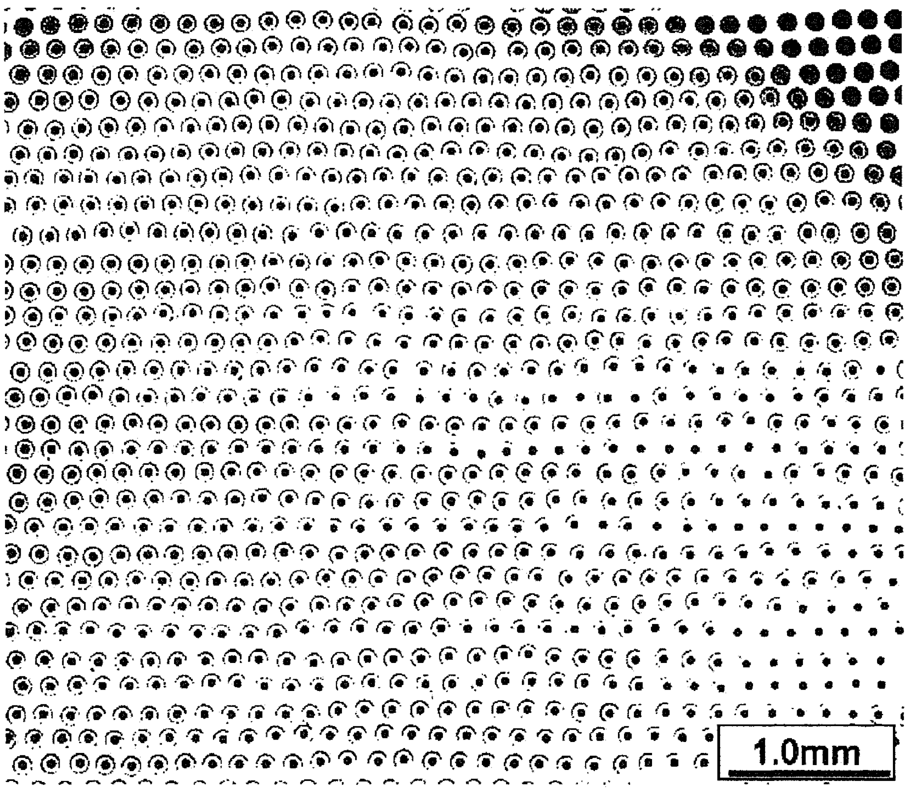

Figure 4. Cross-Section Micrograph of a Ti Matrix Composite. Small, dark centers of fibers are C monofilament substrates for deposition of SiC (light micrograph). Material system: Matrix alloy - Ti-6242; Fiber - SiC about 5.6-mil-diameter SCS-6

Table 2 shows that the strength and stiffness properties of TMCs can exceed those of the Ti matrix (or other $\alpha+\beta$ Ti alloys) by more than a factor of two with no increase in density. So why are TMCs not in widespread use? The foremost reason is cost. Another reason is the availability of SiC fiber for TMCs.

- During the 1990s, when the U.S. government (mainly DoD) was interested in and provided development money for TMC R&D, the most attractive reinforcing fiber was SCS-6™, which was produced exclusively by Textron Specialty Materials (TSM) in Lowell, MA.
- Unfortunately (at least in hindsight), TSM made a business decision not to sell fiber and to instead forward-integrate and sell TMC components or finished TMC mill products. This decision stemmed in part from TSM's negative business experience with development of B-based fiber (Boro-SiC™) for first-generation TMCs.
- Prior to making this business decision, TSM did not develop sufficient TMC-manufacturing expertise to properly position itself as a producer of TMC products with consistent properties.

- A couple of costly and highly visible TMC component test failures using TSM materials called into question the viability of TMCs. In truth, the real issue was one of quality and not the fundamental viability of the TMC material concept.
- These failures led to extreme caution by the government, which promptly imposed stringent quality requirements on TSM. Consequently, TSM became a serious bottleneck for availability of TMCs with acceptable quality. As a result, numerous large, government-funded R&D programs fell behind schedule to the point that some were canceled and others were abandoned.
- Other, less attractive SiC fibers were being developed during this period, including Nicalon in Japan and Sigma in the United Kingdom. Because SiC fiber and TMCs were deemed strategic materials and neither of these alternate fibers was produced in the United States, there was reluctance to experiment with them, and the U.S. government occasionally prohibited their use in federally funded R&D programs.
- The issue of fiber availability became a major distraction from the real business at hand—that is, improving the consistency and reducing the cost of the TMC product.
- In the end, TSM partially relented, but by then such limited interest in TMCs existed that the original opportunity was lost. Furthermore, there still was no concrete evidence that TMCs could be produced with sufficiently consistent properties that they could be considered an engineering material (at any cost).
- In the mid-1990s, a cost study that assumed fiber availability showed that TMCs that meet specification properties could be manufactured for about $500/lb (in 1995 $) if the use volume was about 10,000 lb/year. Finding enough low-risk, high-value applications to consume this quantity of TMCs was not deemed possible.
- Consequently, work on TMCs halted after an investment of about $500 million of U.S. government funds and a (presumably) comparable, but less well known private sector (mostly independent R&D $) investment.
- One lesson from this is that a credible market and cost study should be undertaken before embarking on a major R&D program to develop revolutionary materials such as TMCs.

The foregoing discussion exemplifies the challenges associated with developing and commercializing a revolutionary new materials system. New material concepts originating from nonproduction sources such as national laboratories or research universities should be approached with caution and never be put on the critical path of product design. Even if a new material is vetted through an established production source, the timing of full commercialization should be carefully examined.

Ni-base Alloys

Ni-base alloys have been considered real engineering materials for at least 40 years. Because of their relatively low rate of strength loss with increasing temperature, Ni-base alloys are also called "superalloys." They are commonly used at temperatures well in excess of half of their melting point (T_m), which is the accepted useful limit for most structural metallic materials. There is no other class of structural alloys for which this is true. Ni-base alloys have been used in turbine engines of all types almost since their inception. In fact, Ni-base alloys are a prime enabler of the modern gas turbine engine. As a result, this class of materials has a well-established industrial base for the production of literally all product forms. There is a second class of Ni-base alloys that has excellent resistance to attack in aggressive environments, such as those encountered in petroleum recovery, but these alloys are not of interest for high-temperature applications. To avoid any confusion, the term Ni-base alloys is used in this document to mean *high-temperature* Ni-base alloys.

Ni-base alloys have a high elevated temperature strength because they typically contain Al and Ti, which enable them to be precipitation strengthened. The alloys with the highest temperature capability are strengthened by precipitates of Ni_3Al (called γ'), which have a positive temperature coefficient of strength up to about 1,000 °Celsius. This unusual characteristic accounts for Ni-base alloys' elevated temperature capability. In combination with the ordinary temperature dependence of the matrix strength, the overall temperature dependence is quite low compared with that of any other class of structural alloys. Ni-base alloys also contain alloying additions of refractory elements such as tungsten (W), tantalum (Ta), and niobium (Nb). These additions also improve the elevated temperature strength because they diffuse slowly and strengthen the matrix. Because Ni-base alloys are used at elevated temperatures, they contain Cr additions in significant concentrations to improve their oxidation resistance. The tailoring of the properties of Ni-base alloys has been studied extensively and is quite well understood empirically, albeit less so in a fundamental sense. Current-generation Ni-base alloys contain as many as 10 alloying additions, and interelement interactions make gaining a fundamental understanding of alloying effects a major challenge.

Production of Ni-base alloys has evolved over several decades. Both ingot metallurgy and powder metallurgy methods are used to produce a full range of mill products. New alloys have periodically been introduced in response to specialized needs. Sometimes these alloy development efforts have been so focused on the "specialized need" that retention of other generally accepted Ni-base alloy properties has been overlooked. An example of this is the development of early "low coefficient of thermal expansion" (α) alloys (the so-called low-α alloys). An example of an application for which low-α alloys are attractive is structures in which there are two or more concentric rings connected by radial struts. In the presence of a radial temperature gradient, these alloys tend to crack during thermal cycling situations owing to "thermal fatigue." In principle, the availability of an "Invar type" alloy with structural characteristics reduces the thermal strains and improves the thermal fatigue life. Invar is so named because it uses the magnetic characteristics of Ni to offset normal thermal expansion, resulting in much lower α values. It was recognized that this concept could also be adapted to Ni-base structural alloys. The first-generation low-α alloy, Inco 903, had about half the coefficient of thermal expansion compared with ordinary Ni-base alloys up to the Curie temperature. This remarkable alteration of a physical property was achieved by removing the Cr from the alloy to intensify the magnetic effects of magnetostriction. An unfortunate, unintended consequence of this composition alteration was that the oxidation resistance was seriously degraded. In particular, Inco 903 had a propensity for intergranular cracking in air under stress, rendering it of limited utility as a structural alloy. In fact, the cracking was so severe it was given a name: stress-assisted grain boundary oxidation (SAGBO). A second-generation low-α alloy, Inco 909, contained Si to help improve the oxidation resistance, but this alloy has not gained widespread acceptance, possibly because Inco 903 problems were so severe. Low-α

alloys' superior thermal fatigue resistance gives them great potential for use in TPS; thus, it would be appropriate to revisit this class of alloy after addressing its past problems.

Many of the Ni-base alloys are fusion weldable, but the higher strength grades are prone to cracking without special precautions such as preheating the work piece. Most alloys capable of being processed into sheet also can be processed to create a fine-grained structure. The fine-grained material can be superplastically formed, which is potentially of very important benefit to making light-gauge, load-bearing structures and TPS components. Finally, although almost all Ni-base alloys contain Cr and Al, both of which improve oxidation resistance by forming a stable, protective scale, this scale begins to lose its effectiveness above about 1,000 °Celsius. A number of effective environmental coatings have been developed by the gas turbine industry for Ni-base alloys. The only concern or unknown for application in a reusable SSO vehicle is the durability of these coatings under very high mass flow conditions, such as during reentry.

Refractory Metal Alloys

Refractory metal alloys were discussed briefly earlier in connection with the DynaSoar project. Subsequent to this, there has been limited systematic interest in designing aerospace structures that incorporate refractory metal alloys. There has been interest in using these alloys for high-temperature gas-cooled nuclear reactors, where the operating environment is benign and well controlled, but this has limited relevance to the current discussion of space vehicles required to withstand reentry into the earth's atmosphere. The attraction of refractory metals as a class is the high melting temperature of Mo, Nb, and Ta. Although alloys of these metallic elements exhibit "normal" temperature dependence of mechanical strength, $0.5T_m$ of Nb or Mo is still a higher temperature than Ni-base alloys can withstand. Of the three metals, Nb is by far the most attractive because of its lower density, lower elastic modulus, and better (but not good) oxidation resistance. The clear barrier to use of Nb alloys is their reactivity in air at high temperatures. There are coatings for these alloys, but at 1,250 °Celsius, a breach in the coating will cause immediate, catastrophic failure.

A reasonable design practice for critical structure, such as the TPS, is to disallow use of any coated material if the material fails catastrophically when the coating is breached. If this practice were the norm, then refractory metal alloys would not be usable in the TPS or other critical hot structure in a reusable SSO vehicle.

Ceramic Matrix Composites

Ceramic matrix composites (CMCs) are attractive because they are much tougher than, but retain the high-temperature capability of, monolithic ceramics. CMCs' toughness is derived from the fibers used to reinforce the ceramic matrix. When a CMC is loaded in tension and the stress in the ceramic matrix reaches a critical value, microcracks develop. This stress level is known as the matrix microcrack stress. In a monolithic ceramic such as silicon carbide (SiC) or silicon nitride (Si_3N_4), microcracks would propagate rapidly, causing immediate failure. For this reason, monolithic ceramics are not suitable for tension-loaded applications. This characteristic is mitigated in CMCs by reinforcing the ceramic matrix with high-strength ceramic fibers, typically Al_2O_3 or SiC, whose role is to bridge and arrest the cracks. Under continued loading, the fibers also will eventually fracture at their weakest points, which, statistically, is often somewhere in the matrix rather that at the mouth of the crack. These fractured fibers then must be pulled out of the matrix during crack extension under continued loading, an action that dissipates additional energy. Therefore, the total energy expended during fracture of a CMC is much greater than it would be in the case of a monolithic specimen of the ceramic matrix. The crack bridging and pullout mechanisms of toughening, with the attendant increase in energy dissipation during fracture, are the conceptual foundation for CMCs. In SiC fiber-reinforced SiC matrix (called SiC-SiC) CMCs, the toughness can be as high as 10 times that of the unreinforced matrix. Clearly, the nature of the fiber-matrix interface determines the resistance to fiber pullout after fracture and controls the toughness. In high-temperature CMCs, the fibers are coated to achieve an intermediate fiber-matrix interfacial bond strength that optimizes pullout toughening. If the interface is too weak, the fibers pull out too easily. If it is too strong, the fibers break without any pullout. Neither of these situations maximizes the toughness. Therefore, the challenge in creating a tough CMC is not only creating this interface but also finding a coating that will remain stable over time during exposure to elevated temperatures. As service temperature increases, this latter requirement becomes more challenging. For extended service, SiC-SiC CMCs currently are limited to about 1,400 °Celsius, but this is higher than the capability of metallic materials.

The matrix microcrack stress is also significant because once the matrix develops microcracks, it allows the environment to gain access to the fiber-matrix interfaces. The role of environment can be to alter the nature of the interface and reduce the extent of pullout toughening. The obvious difficulty with this is that the properties are better initially than after a period of exposure in service. This leads to a nonconservative design that is dangerous. Absent an arbitrary knockdown, there currently exists no means of estimating the reduced toughness as a function of service life.

Nevertheless, CMCs are the material class that holds the greatest promise of defeating the temperature limits of current metals. An important issue at present is the limited industrial base for producing CMCs and the even more restrictive range of suitable ceramic fibers available for use as the reinforcement in CMCs. The cost of CMCs is currently very high but should come down with increased demand, as happened with TMCs. Even so, CMCs will always be expensive, making development of efficient designs that make optimal use of this class of material important.

Carbon-Carbon Composites

Carbon-carbon composites (C-CCs) consist of carbon fibers in a typically amorphous carbon matrix. In principle, this class of materials is very attractive because it has very low density and high strength at elevated temperatures. C-CCs also have good oxidation resistance at very high temperatures (1,200 °Celsius) because they use an Si-rich coating that forms a stable, protective SiO_2 film on the surface. This coating system is "self healing" if breached because, among other things, it is glassy (viscous) and reforms over cracks. The principal vulnerability occurs at intermediate temperatures where the viscosity of the glass is high enough that it may not flow rapidly enough to heal a breach, exposing the reactive C substrate to the damaging oxidizing environment. As with CMCs, there also are some questions concerning long-term changes in the fiber-matrix interfaces that are at least partially responsible for the loss of impact toughness over time, as experienced in the Columbia accident.

Therefore, instabilities that cause time-dependent reductions in toughness must be understood before these materials can be considered candidates for future TPS applications.

Titanium Aluminides

Titanium aluminides are intermetallic compounds that form between Ti and Al. There are three such compounds—Ti_3Al, TiAl and Al_3Ti—but the one of principal interest is TiAl, often referred to as gamma titanium aluminide or γ-TiAl. This compound contains about 35 weight % Al, but ternary and quaternary alloys based on TiAl typically contain somewhat less than this amount. γ-TiAl is interesting for intermediate-temperature applications for several reasons. These include a lower density, at least a 125 °Celsius higher temperature capability, and better surface stability (oxidation resistance) compared with conventional Ti alloys. The surface scale that forms during oxidation of γ-TiAl is Al rich and is more protective than the TiO_2 scale that forms on conventional Ti alloys. Like most other intermetallic compounds, γ-TiAl has very limited ductility at temperatures up to about 600 °Celsius but is not brittle in the classical sense; that is, it has the capability to deform plastically before fracturing. The ductility increases as temperatures increase and is quite good at normal service temperatures.

γ-TiAl alloys' limited ductility poses manufacturing problems for any wrought products. These issues are manageable but reduce product yields, which in turn impacts cost. Recent work in Austria has demonstrated that some γ-TiAl alloys can be made into sheet. This is particularly relevant to applications such as metallic TPS. Other work has demonstrated that some γ-TiAl alloys can be conventionally forged, albeit with care. Should one or more of the efforts currently under way to produce low-cost, prealloyed Ti powder prove successful, the availability of affordable γ-TiAl alloy powder could be a major cost breakthrough.

The development and maturation of γ-TiAl alloys has taken more than 30 years. This is in part due to the lack of actual production applications because of designer's concerns regarding the limited low-temperature ductility of all the γ-TiAl alloys. Today, finally, there are several applications of alloys based on γ-TiAl, the most significant of which are two stages of low-pressure turbine blades in the engine General Electric is providing for the Boeing 787 and the growth 747 (called the 747-8). These blades are cast to near net shape and will operate at temperatures up to about 750 °Celsius. Another application is a cast γ-TiAl turbocharger rotor Mitsubishi uses in a production auto because it has half the mass of an equivalent Ni-base alloy component.

A potentially ideal, high-volume application of γ-TiAl alloys is exhaust valves for autos. If the low-cost powder mentioned earlier becomes a reality, this factor, coupled with the current pressure to improve fuel consumption in American cars, could make this prospect a reality. The point of this short discussion of non-space-related applications of γ-TiAl alloys is that realization of any or all of these applications will create a stronger production base for this class of alloy, which in turn will make their introduction into future space platforms easier and more cost effective.

PROPULSION SYSTEMS

Among propulsion systems, only reusable rocket engines are considered here because they hold potential for significant progress that could reduce the cost of placing payloads into orbit.

The basic concept of the space shuttle main engine is still viable, but the durability of the materials used to make the hardware has been a major expense and source of concern for NASA. Much of the concern is related to the effects of hydrogen on the Ni-base alloys used in the turbo pumps. The turbo pumps are the heart of a liquid-fueled hydrogen-oxygen rocket engine because they deliver the fuel and oxidizer to the thrust chamber. The turbo pumps rotate at a very high speed—up to 35,000 rpm. Such speeds create enormous centrifugal stresses in the rotating components, particularly the disks that hold the air foils used to extract work from the hot gas stream. One side of the turbo pump rotor operates in the hot gas stream created by the combustion of the oxygen-hydrogen mixture; the other end is a cryopump that operates at cryogenic temperatures, either in liquid oxygen (90 K) or in liquid hydrogen (20 K). The hot side, a turbine, resembles the turbine rotor in a gas turbine engine and is made of many of the same Ni-base alloys. The cryopump is made of Ti alloys.

The turbine essentially operates in a hydrogen-rich supercritical steam environment at a maximum temperature of about 1,050 °Celsius. The turbine disk is a forged Ni-base alloy, and the air foils are single-crystal investment castings. Neither material is well suited to operate in a hydrogen-rich environment; however, no other material class can withstand the operating temperatures and has better hydrogen tolerance—a situation that still prevails today. The mechanism of hydrogen-induced cracking is much better understood today as the result of extensive research over the past 25 years. Therefore, it would be very useful to use this improved understanding to design a Ni-base alloy that has improved hydrogen tolerance. If successful, a turbine that has improved resistance to hydrogen cracking would greatly reduce the intermission refurbishment time and cost.

Traditionally, the cryopump rotating parts have been made from forgings of the near-α-phase Ti alloy Ti-5Al-2.5Sn (Ti-5-2.5). This alloy has been chosen for its superior notched tensile strength compared with Ti-6-4 when tested at cryogenic temperatures. Hindsight suggests it is unclear that notched tensile strength is the best criterion for selecting a cryogenic rotor material. The use of notched tensile strength originated with steels. Here, the hydrostatic stress state at the notch root could trigger the onset of brittle fracture of the type seen in smooth tensile tests below the ductile brittle transition temperature (DBTT). Ti alloys do not exhibit a DBTT; therefore, true fracture toughness measured at the relevant temperature is a more accurate indication of the fracture resistance of the rotor. An examination of the limited available fracture toughness data for the two alloys shows no clear advantage in using Ti-5-2.5. Because Ti-5-2.5 is more difficult to produce than Ti-6-4, using Ti-5-2.5 adds cost to the cryopump. Other, newer, higher strength Ti alloys—for example, several of the newer β-Ti alloys—may be even better suited than Ti-6-4 is for this application. The critical question of how much toughness is really required cannot be answered at present because the notched tensile ratio has been the deciding criterion. Therefore, a new design of a cryopump potentially can benefit from these new materials and from abandoning the use of notched strength as a design paradigm.

The combustion chamber of a rocket engine is made from a copper alloy and is cooled by passing liquid hydrogen through channels in the outer wall. The space shuttle main engine combustion chamber is made of NARloy-Z, a Cu-3%Ag-0.5%Zr alloy that has excellent thermal conductivity and better strength than the commercially available Cu-0.15%Zr alloy known as AMZIRC. While a lighter combustion chamber would be welcome, developing hydrogen-resistant turbine materials would be a much better use of available rocket engine alloy development resources.

SUMMARY AND RECOMMENDATIONS

This document has attempted to review a range of materials that may have promise for all aspects of aerospace platforms. Because of the breadth of this topic, discussion has focused on opportunities and associated risks, but with little technical detail. Practical issues such as manufacturing capabilities and costs and the availability of materials also have been addressed where appropriate. Another recurring issue is the ability to achieve better design efficiency through a design synthesis process that concurrently treats form, fit and function, manufacturing capability, and materials capability as equally important design constraints. Although this approach has yet to be successfully used for a major high-performance structure project, engineering is nearing a state of development where it may now be feasible.

If the space shuttle is used as a benchmark, it is clear that numerous opportunities will exist to improve structural efficiency. This possibility reflects the progress made since the shuttle and its engine were designed. Improved durability and reduced operating cost also are possible. The issue of long lead time holding back these improved design methods is most often one of timely availability of new, attractive materials that are mature enough to be used in a system without increasing the risk of failure. The materials community is actively developing and using computational models and simulation methods to address this concern. Despite some progress, further advances are required to bring computational materials engineering to the desired level of maturity.

1. Hatch, J. ed. Aluminum, Properties and Physical Metallurgy. American Society for Metals, 1984.

2. Chawla K. Composite Materials, Science and Engineering 2nd edition. Springer, 1998.

3. Baker, A., Dutton, S. and Kelly, D. Composite Materials for Aircraft Structures 2nd edition. AIAA, 2002.

4. Lütjering, G and Williams, J. Titanium, 2nd edition. Springer, 2007.

5. Sims, C. and Hagel, W. The Superalloys. Wiley, 1972.

6. Mishra, R. and Ma, Z. Friction Stir Welding and Processing, Materials Science and Engineering. v.50, pp. 1-78, 2005.