State-Of-The-Art Advanced Composites
And Carbon/Carbon Technology

Randy Lee

April 2000

All concepts proposed and developed here are entirely original and are the analysis, perceptions and opinions of Randy Lee

This essay draws heavily from experiences obtained during my tenure at LTV Aerospace & Defense but also utilizes activities and concepts acquired since that time in other high technology areas. In order to minimize redundancy, the narrative portion of my resume should serve as the background and introduction to the following discussion.

Reinforced Carbon/Carbon (RCC) has been used on the Space Shuttle since the late 1970's with great success, and its predecessor was employed in the Apollo program shortly after its accidental discovery at LTV in the 1960's. In more recent years, the carbonized rayon fiber-based RCC has not attracted quite the attention that polyacrylonitrile (PAN) carbon fiber composites have. Carbonized PAN reinforcements offer tensile strengths over 5 times greater than their carbonized rayon counterparts, but they are also much stiffer with elastic modulus values 5 or 6 times greater. It goes without saying that, due to the final 3000° coating process and expected service environments, thermal expansion (CTE) differences between densified carbon/carbon substrates of either form and the surface-fused SiC coating are very important. Coating-to-substrate CTE problems associated with the Space Shuttle's RCC have been quite manageable, for the most part. However, the same cannot be said of more recent coated PAN-based Advanced Carbon/Carbon (ACC). One thing is certain, the coating process always degrades the substrate's overall mechanical properties significantly regardless of which substrate material is used but the effects are much more devastating with PAN reinforced composites.

Expansion properties for both fully densified RCC and ACC substrates are quite similar, but modulus (stiffness) characteristics differ by such a wide margin that significant cracking and failure of the coating on ACC surfaces has often been the result after thermal (plasma) testing or, quite frequently, during cool down of the coating process itself. Longitudinal CTE values for rayon and PAN carbon fibers are close to zero while monolithic β-SiC is around 4 (both carbon fiber forms are currently produced by Amoco Performance Products). While crystalline planes are generally oriented along the longitudinal dimension, expansion occurs within the 'd' spacing (across the fiber diameter). Additionally, the tensile modulus for carbonized rayon fibers runs about 6 (MPSI) and that of PAN carbon fibers around 35 (six times greater!). Thus, rayon fibers should have a higher tendency to 'move' with the SiC coating during both processing and atmospheric re-entry episodes. It is now a historical fact that substrate-to-coating CTE mismatch is a showstopper for ACC materials and is not only manageable for RCC, but is essentially unimportant. Because of its higher strength properties, PAN-based ACC carbon/carbon was selected by NASA and other industries for continued development. Since that time however, coated ACC has had a long, difficult history meeting expectations for processability and attempting to compete with the longevity triumphs of RCC. While RCC had proven itself very well for certain space plane skin panels, the demand was growing for a producible carbon/carbon form that could be used in airframe structural applications.

During the 1980's, much of the carbon/carbon activity at LTV and throughout the industry was focused on addressing developmental problems with ACC PAN-based material in two major areas: thermal/mechanical incompatibilities associated with the coating phase (as previously mentioned) and poor interlaminar adhesion often resulting in ply-to-ply delaminations. RCC has also exhibited similar effects over the years including coating microcracks as well as occasional delamination problems but the degree of severity is quite insignificant when compared to ACC materials. While carbonized PAN fibers provide enormous strength benefits for ACC products throughout the laminar plane, interlaminar association between the plies is considerably weaker. My team undertook the task to develop and refine a specialized test method for measuring ply-to-ply adhesion and this resulted in a modifed version of the 'flatwise tension' or interlaminar tensile test. Interlaminar tensile (ILT) ranges of 800-1200psi were typical for densified RCC laminated substrates (prior to coating) while densified ACC substrate panels usually failed around 400-800psi (I personally directed over a thousand tests verifying these averages). After 3000°F SiC pack mix conversion coating, ILT values would drop anywhere from 40% to 70% with RCC exhibiting the lowest strength degradation (near the 40% end) and ACC the highest (60-70% all the time).

There are at least two reasons that accounted for these results. One is directly related to the surface and bulk characteristics of the fibers and the other deals specifically with the matrix phase. Examination of fiber surface morphologies reveals that PAN carbon fiber bundles are smooth, slick and relatively straight while rayon bundle surfaces are irregular (crinkled) and much more porous (verified by Amoco). This is due to a ‘looser’ molecular structure for carbonized rayon and a more organized (and graphitic) arrangement for PAN fibers. The net result is an increased ‘nesting’ effect between fabric planes within rayon-based RCC laminates and a high degree of association between adjacent fiber bundles. Additionally, the higher porosity of rayon carbon fibers serves to enhance matrix permeation into the fiber itself allowing the composite phases to become more intimate with one another; rayon fibers are more easily ‘wetted out’ by the resin phase while PAN fibers almost tend to repel the matrix. These properties also help to explain why PAN carbon fibers are much better heat conductors than rayon. Secondly, viscous phenolic matrix densification resins are used for ACC while RCC has always utilized furfuryl alcohol resin mixtures for substrate densification (in addition to furfuryl alcohol, this mixture also contains formaldehyde and possibly phenol). While traditional phenolic resins are somewhat safer and easier to handle than furfuryl alcohol mixtures, there are strong indications that they do not permeate the fiber porosity, matrix porosity, bundle interstituals and voids near as well as furfural alcohol. The low viscosity and lower surface tension of furfuryl alcohol significantly enhances penetration and wettability effects when compared to the highly thixotropic phenolic resins.

While the ultimate objective is a high strength oxidation resistant product (i.e…ceramic coated carbon/carbon structure), the initial molding concept and laminate fabrication process play a key role in how well the substrate responds to subsequent densification steps and the coating phase. A common starting material might consist of one of Fiberite's carbonized PAN fabric/phenolic resin prepregs. Hundreds of experiments and fabrication trials were performed to determine the optimum lay-up and bagging configurations, B-staging conditions, autoclave parameters, time-temperature profiles, bleeding/breathing requirements and the effects of post-curing. Typical fiber volumes ran around 60-70%, resin contents in the 25-35% range, bulk densities 1.6-1.7 and apparent porosities of 3-8% (measured by fluid intrusion or ASTM water boil). Again, I have personally directed and perform thousands of tests to verify these ranges. Phenolic resins give off condensation products, mainly water, creating pores so they must always be cured under pressure regardless of the specific application (composite fabrication and substrate impregnation).

After the initial pyrolysis (or carbonization) step at 1500°F, substrate density drops to 1.3-1.4 and the apparent porosity increases to 20-25%. This material is highly porous and delicate – essentially useless for any kind of structural application. However, successive cycles of densification carried out by vacuum/pressure resin imprenation followed by pyrolysis (matrix carbonization) bring the substrate properties back up to significant ranges. After four densification cycles, ACC laminates have densities, porosites and matrix fractions almost identical to that of the molded composite, except the matrix is now entirely glassy carbon. The fiber volume remains constant throughout the entire substrate process (as it is established in the composite design and fabrication stages). ACC substrates have commonly received four densification cycles prior to coating (designated as ACC-4). The properties are somewhat different for RCC systems which are based on the lower density rayon carbon fiber reinforcement and the more penetrating furfuryl alcohol densification resin. Three cycles are typical for producing an optimal RCC-3 substrate which effectively accommodates the final conversion coating.

Simple relationships were formulated to help define and explain specific principles involved as the materials underwent processing. This comprehensive model was formulated independently by the author and is covered in a separate report. Eventually, critical substrate properties (density, porosity, component fractions) were correlated with weight gains from impregnation, pyrolysis weight losses and net carbon matrix weight gain after each completed cycle. At any given processing (carbon) state, these properties can be approximated from a common but modified expression often starting with only simple measurements of the panel weight and dimensions (this applies to the initially cured panel and any subsequent densification state):

$$\rho_b = \left(f_w \rho_f^{\;-1} + m_w \rho_m^{\;-1} \right)^{-1} (1 - p)$$

where ρ_b is the panel's bulk (or geometric) density, f_w the fiber weight fraction (fiber volume is constant but fiber weight varies), ρ_f the fiber density (which is a constant but difficult to measure and runs around 1.8-1.9 for heat treated PAN fibers), m_w the matrix weight fraction (determined from the initial resin content), ρ_m the matrix density (pressured cured phenolic resin is ~ 1.24 and ambient carbonized resin ~1.42 as measured numerous times), and p which is the apparent porosity of the substrate at the state in question.

Most likely, the theoretical porosity is much higher than that measured by fluid intrusion or impregnation. Mercury porosimetry was found to be inadequate (possibly damaging the sample) while ASTM water boil or modified water impregnation (a method developed by myself) generate porosity volume fractions which correspond very closely with the pore volume occupied by resin after impregnation cycles. Obviously, with CVD/CVI densification techniques, this approach would be less than adequate (corresponding more closely with He picnometry measurements). However, with the resin impregnation densification method, glassy carbon is deposited within the composite as the matrix is densified by the succession of carbonization cycles.

The matrix weight fraction at any carbon state i within the process can be expressed by the following generalized product:

$$m_{wi} = 1 - (1 - r_c)(1 - n_{l0})^{-1} \prod_{1}^{i} (1 + n_{g(i-1)})^{-1} (1 - n_{li})^{-1}$$

where r_c is the initial resin content of the laminate (matrix weight fraction at the molded state), n_l the fractional panel weight loss after pyrolysis and n_g the fractional panel weight gain after resin impregnation and cure. Once the 'as-molded' resin content is known, the matrix weight fraction (or matrix content) can be estimated throughout the densification process by simply measuring the weight of the part before and after each step. The expression is not as complicated as it looks: the matrix content after the first pyrolysis is simply $1 - (1 - r_c)(1 - n_{l0})^{-1}$, and each successive densification cycle modifies the product on the right side by incorporating the corresponding weight gains and weight losses. Eventually, it was found that each of the increasing parameters (primarily density and matrix content) could be approximately traced over the entire densification process by a simple increasing function of the order,

$$P = A\left(1 - e^{-ki}\right) + C$$

This expression also describes most substrate mechanical properties as well (flexure and ILT were successfully simulated). The primary decreasing variable, open porosity, was determined to follow the exact opposite of this form, a simple Arrhenius equation: $Ae^{-ki} + C$. Of course a much more extensive system of process/property concepts and formal definitions were developed than presented here as published in numerous reports over several years (including definitions of composite condiitons in the 'bi-matrix' state). These important statements set the stage for systematic characterizations of the carbon/carbon process in an effort to quantitatively define the most suitable ACC substrate to undergo the final SiC surface conversion operation successfully. Ultimately, the goal was to link fully densified substrate properties to the initial molding parameters so that optimal substrate could be conceived of during the design stages rather than the tedious tweaking and step-to-step troubleshooting that was typical of almost every article made. Unfortunately, defense cutbacks and program reductions, reminiscent of the late 1980's, hampered the final completion and advancement of these concepts. The Carbon/Carbon Technologies group at LTV, once 70-80 people strong, has long since vanished. This work has not yet been completed . . .

During that period, a number of alternative concepts were also explored. In addition to the more common materials used for producing ACC and RCC, my team investigated several non-traditional fabric forms (at the time), techniques for enhancing interlaminar strength and alternative molding methods. Numerous samples of Nextel fabric, braided PAN and preformed weaves were subjected to impregnation, casting, pressure curing and modified vacuum bag autoclave curing in order to study their behavior during the densification process. Even though 3 dimensional forms are not laminates, ILT values for most of these substrates were well below the averages associated with standard ACC laminates. However, advanced concepts applied to special 3-D configurations are definitely feasible. Even though rare to acquire at the time, crenulated PAN fabrics performed much better than the standard 8 harness weaves and are still recommended for further development (analogous to but still not as good as RCC). A NASA sponsored project to study Z-stitched, debulked lay-ups was heavily investigated but the results were predictable from start – unless the fabric is dry, this approach produces cured laminates that fail right along the stitch line. The particular pitch fabric samples studied were not impressive however, this reinforcement candidate definitely deserves further investigation. Concepts utilizing nanotube carbon fibers look quite promising today but were never studied during this era (they were not even available back then!). While the actual time and efforts used during the investigations of these alternative concepts was much greater than implied here with this short summary, the discussion will focus on the principles that received the greatest amount of attention specifically related to more heavily studied carbon/carbon materials and manufacturing processes.

In general, carbon/carbon densification consists of matrix densification by deposition of carbon within the porosity of the composite. For the resin impregnation approach typical of LTV's RCC and ACC programs, densification materials have utilized high carbon thermosetting resins producing maximum char yields upon pyrolysis. This limits the field to just a few selected classes, primarily phenol-formaldehyde and furfuryl alcohol mixtures.

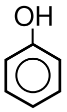
Phenol
(ACC)

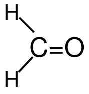
Formaldehyde
(ACC & RCC)

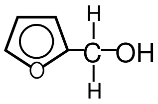
Furfuryl Alcohol
(RCC)

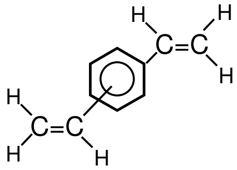
Divinyl Benzene
(Developmental)

Rather than graphically going through all the polymerization schemes and specific reaction mechanisms involved here (these are explicitly covered in a separate report), the process is briefly described and then concepts deemed more relevant to the densification phase are discussed. The resin manufacturer generally reacts and processes the resin constituents through one or two stages. Stage A comprizes the initial ionic reactions between monomers to form hydroxylated oligomers or 'pre-polymers' which are identified as either 'resoles' or 'novalacs' depending on whether acid or base catalysts are used. For the phenolics, phenol and formaldehyde are reacted to form dimers and trimers of methylene linked methylol phenols. Furfuryl alcohol reacts with itself and then formaldehyde to form dimers and trimers of methylene linked methylol furans. The vendor then adds additional catalysts (usually an acid) and more of one or both monomers and possibly carries the product partially into the Stage B phase or to a point specified by company guidelines or by the customer. The customer receives the resin near the start of the B phase and then completes the Stage C upon curing in a process often referred to as 'crosslinking' (the customer generally adds an acid catalyst to furfuryl alcohol just before use). For phenolic and furfuryl resins, polymerization and crosslinking are indistiguishable since reaction advancement occurs near isotropically. Both resins are classified as 'condensation' polymers, so the by-product of water is released throughout all three stages and, after incorporation of resin into end product material phases (i.e., substrate or fabric impregnation), this water must carefully be monitored and allowed to expel properly. The fully cured form of both phenolic and furfuryl resins represents a vast network of aromatic rings joined by methylene (carbon) groups. Carbon links between aromatic groups throughout the backbone inhibit structural degradation of the polymer during pyrolysis which would ordinarily occur with most other resins (epoxies, urethanes, imides, esters, acrylates, etc…) because oxygen or nitrogen groups comprize the major backbone links.

Phenolic and furfuryl alcohol resins must be cured under pressure because the sporatic release of solvents, monomers, partial products and moisture causes significant bubble formation which reduces the remnant resin solids, creates massive porosity and voids and an annoying mess to be cleaned up afterwards. Some of my investigations included imprenations cured at ambient pressure indicating that efficiencies (weight gain) are reduced by 40-60% where the resin loss is deposited inside the curing chamber and on the outside of the part (rather than within the substrate's porosity). Divinyl benzene is proposed here as a possible alternative thermosetting polymer for generating very high carbon yields and one which cures by a different mechanism. Free radical near-isotropic polymerization of this monomer could possibly be accomplished at reduced pressure, minimal heat application and without the formation of water or residual OH groups (problems associated with residual hydroxy groups are identified below). While availability and safety hazards associated with this product appear to be less than desired, total cycle time reductions as well as notable overall cost reductions would be quite possible.

The majority of carbon/carbon produced at LTV in the 1980's utilized phenolic and furfuryl alcohol resins throughout. Both ACC and RCC prepregs came impregnated with phenolic resin but ACC underwent phenolic resin densification while RCC has been furfuryl alcohol densified since its inception. Extensive curing studies have strongly indicated that a moderate cure ramp (3-5°F/min) with 1 hour holds at 180°F and 325°F produces the best (optimal) product and is considered to be the standard cure for phenolic resin systems. However, at the end of a typical pressurized curing cycle to 300-325°F, the cured product contains latent properties often overlooked which usually lead to weak planes and seemingly unexplainable delamination problems downline.

After standard curing, a certain level of residual, unreacted OH groups are still present and trapped within the cured phenolic matrix structure. These hydroxy entities are nothing more than the original OH groups attached to phenol but have become shielded during the polymerization stages and now require even greater energy for expulsion from the cured material. Likewise, the cured furfuryl network still contains the original cyclic ether group which requires greater energy for bond degradation permiting release of the oxygen in some form or another. Ideally, these oxygen links should be removed with classical carbon-carbon bonds being formed in their place to adjacent sites or molecules. We could possibly refer to this process as 'post-crosslinking' since a higher temperature follow-up cure or 'post-cure' is the most likely remedy. Indeed, repeated trials and tests have shown that a 400-425°F post-cure after each standard cure has completely eliminated delamination problems and enhanced the mechanical integrity of the final substrate. A two phase curing process consisting of the 325°F cure (*and cool down*) followed by a 425°F postcure has been shown to produce superior parts when compared to a single cure to 325°F or 425°F. Cure cycle studies indicated that the resin retains its liquid properties throughout the heating cycle and does not actually begin to 'harden' until the system enters the cool down phase when the intrinsic thermal energy begins to convert into mechanical stresses within the rapidly solidifying matrix. It is likely that the 425°F postcure enhances structural relaxation throughout the solid matrix and facilitates residual oxygen/water removal where a pyrolysis cycle immediately following the 325°F cure rapidly forces these constituents out possibly damaging the matrix structure. The postcure cycle (*and cool down*) helps to ease out ‘pyrolysis volatiles’ from the matrix maximizing the fractional carbon content and minimizing the level of thermal clevage points (matrix discontinuities) within the final carbonized product.

Another area of concern is the up-ramp or rate of applied temperature increase during the cure profile to a laminate or impregnated body. Essentially all resins used throughout industry are 'solutions' containing, at the very least, a mixture of reacting monomers, dimers, trimers, oligomers, catalysts and solvents, and many products also incorporate accelerators, surfactants, lubricants, inhibitors, etc… When the applied temperature increases, the motion of each component increases as it absorbs energy and tends to migrate toward the periphery of the laminate. For pressurized autoclave cures utilizing bagged laminates under vacuum, peripheral migration is extremely pronounced. Thus, the components begin to separate and travel at a different speeds along the fiber surfaces and so the process of polymerization competes with the migration effect. This has been observed numerous times during cure cycle studies by halting the cure process at various points along the profile and examining the laminate and bleeder cloth materials. 'Differential migration' of resin constituents during cure might be analogous to the principle at work inside the separation columns of chromatography equipment. As one might expect, solvents are the first entities to leave the system, followed by monomers (phenol, formaldehyde) and lower weight pre-products. Water is expelled throughout the entire curing process as it is formed from condensation reactions.

For larger panels, periphery areas closer to the edges permit less restrained resin flow than center regions where components often can only move across the laminate planes (the Z direction). Hence, many panels result in 'resin rich' central regions and 'resin starved' outer regions. Also, since bleeder material typically interfaces the panel on the upper side with the mold tool surface on the opposite side, panels will generally have higher resin contents near the mold face of the part. The thicker the laminate, the greater the effect will be. At some point (of thickness), the balance of resin content across the thickness becomes excessively large and quite difficult to control using conventional lay-up assembly practices. Attempts were made to explore the effects of faster and slower curing (ramp) rates which eventually lead to the optimum, standard cure cycle mentioned earlier. Results of accelerated ramps were usually unfavorable since rapid curing of thermoset resins within a large prepreg or impregnated network generally produce wide variations or gradients in ‘degree of cure’ throughout the part. Slower curing rates increase the 'chromatographic' effect and produced regions of low resin content and areas often with no visible resin at all. This might be analogous to the effects observed when the mixing ratio of multi part epoxy resin systems is altered resulting in incomplete curing and poor reaction products. Since the resins used for most composite systems are mixtures of the reacting components, differential migration of resin constituents during cure is associated with essentially all resin systems. The often used terms of ‘resin rich’ and ‘resin starved’ may be somewhat over simplistic since the problem is related more to unbalanced constituent ratios than simply high or low resin content. While slight variations in the cure cycle may tend to accommodate different panel thicknesses and geometries, it goes without saying that optimization for any given resin system is a requirement and is often determined empirically. Most of these principles also apply to simple FRP, graphite/epoxy and polyimide/BMI systems as well, but since these systems are not ordinarily subjected to extreme temperatures during processing, the effects often remain latent, sometimes resulting in unexplainable problems downline, later in the field or they might never be noticed at all. Needless to say, a firm understanding of advanced composites can facilitate marked improvements in product quality for the more traditional composite material systems.

In general, the overwhelming majority of thermoset polymer-based composites will tend to exhibit some or all of the previously mentioned effects to varying degrees, even if the consequences are not apparent from an outside perspective. Resin distribution variations, weak interlaminar planes and geographical degree-of-cure variations are latent problems associated with all polymer resin fabric/fiber systems regardless of the particular resin, type of fiber or fabrication method employed. They have to be dealt with sooner or later and much sooner when it comes to carbon/carbon composites. Quite frankly, the experiences gained from carbon/carbon manufacturing technology provide an excellent learning tool for the simpler composite systems considering the extensive level of post-fabrication processes this material is subjected to (repeated high temperatures, pressure, vacuum, infusion, etc..). Control over the original design and molding/fabrication phase must be maintained and well understood, more so than with typical polymer matrix systems which are fabricated and then submitted to their respective assembly and field applications after relatively low temperature and oxidation exposures (coated RCC endures >2000°F almost indefinitely). While many composite specialists may consider some of the more simpler resin/fiber systems to be 'advanced' (such as simple graphite/epoxy), carbon/carbon truely is an advanced composite material.

Laminated carbon/carbon technology has taught us that the best resin-based composites are those that are subjected to post-curing operations after the initial standard cure and are fabricated from prepreg sheets that have been 'B-staged' prior to laminate lay-up or debulking. Generally, B-staging consists of subjecting the prepreg (rolled out or cut into sheets) to slightly elevated temperature (in an air-circulating oven perhaps) in order to increase resin viscosity and advance polymerization. Some prepreg producers may perform staging operations in-house but the optimum temperature and hold time are parameters the composite manufacturers will often want determine themselves for their particular application. Some end user fabricators (such as Lockheed, Bell and apparently many of the major aerospace firms) have performed this operation on layed-up laminates during debulking phases. This is incorrect since resin component separation significantly increases with every degree rise in temperature. Clearly, an assembled laminate under full vacuum with a little heat applied (but still well below the curing range) is undergoing extensive resin migration and separation. This treatment actually competes with the true intent of staging and counteracts its intended benefits. The longer the holding period during this frivolous staging approach, the more separated the resin becomes and the poorer the laminate that results. It would make better sense to quickly bring the green lay-up into curing temperature ranges and complete the cure process it than to 'stage' it in this manner. These fabricators should let the prepreg manufacturer staged their raw material for them. Properly performed, B-staging is carried out on single layer, unstrained prepreg so the components may react *in place* minimizing resin flow and separation across the fiber surfaces. Effective B-staging of free standing prepreg slightly advances polymerization (in place), increases resin viscosity and consumes more of the mobile reactants so that during the final curing phase, component ratios are more evenly distributed throughout the laminate and the thicker inner regions are not so 'resin-starved'. Indeed, this is the very purpose for staging in the first place.

For carbon/carbon systems, the molded and postcured article is subjected to pyrolytic conversion at ~1500°F over a specific time-temperature pyrolysis cycle thus carbonizing the cured resin and forming a highly porous composite system containing an amorphous, inorganic carbon matrix. Densification consists of successive cycles of resin impregnation/cure and pyrolysis which gradually fill the open porosity. About the same basic curing cycle that is used to mold the initial phenolic resin/fabric laminate is also employed during the impregnation densification phases as well. However, special attention has to be paid to the particular processing techniques employed during densification phases to minimize the formation of ‘closed porosity’. For instance, the use of ‘double impregnations’ after the initial pyrolysis was practiced by LTV for many years but there is evidence that this attempt to accelerate the resin impregnation/densification process produces more problems than it is worth. Specific studies performed by myself strongly indicated that a certain fraction of the pore/void network was actually being closed off by the second impregnation and attributed to subsequent part failures (delaminations) downline. While a postcure after the first impregnation/cure tends to alleviate the problem, double impregations without the intermediate pyrolysis step should be avoided. Expansion of the gases during the pyrolysis process following a double impregnation step might fail the laminate at that point or more often will only weaken the interlaminar network and not actually produce failure until some later step (sometimes as far down as the coating process or later on in the field).

The ultimate goal is to create and maintain an open and continuous pore/void network throughout the composite and throughout the entire densification process. As the sequence of densification cycles progresses, these ‘pore tunnels’ should be decreasing in length (and width) as carbon is deposited within the inner regions first and works its way toward the outer volumes. Ideally, the last carbonization cycle is forming matrix deposits exclusively within smaller pore volumes near the periphery of the part, for if resin is still penetrating internal regions at that point in the densification process, then the earlier impregnation/cure/pyrolysis cycles were in effective. Panel weight gains from consecutive resin impregnation/cure/pyrolysis cycles have been proven to follow the increasing exponential forms described eariler when plotted at each bimatrix state. The same can be said of the final carbon gain from each cycle as plotted at each carbon (pyrolysis) state. Fluid measured porosity fractions have been used to approximate the expected weight gains from impregnation/cure; i.e…one can easily estimate the expected resin pick-up by knowing the apparent substrate porosity beforehand. However, it is quite feasible that extrapolation of certain equations developed during these studies (some were presented earlier), could permit the estimation of resin weight gains and carbon gains throughout the entire densification process simply by measuring the initial resin content and the panel weight loss from the first pyrolysis. These concepts need further development and refinement. The possibility now exists to physically link composite design parameters to the coating phase.

Solids content for typical phenolic resins is about 65-70% (when cured to 325°F under 80-100psi). Post cure (which removes residual but critcal volatiles) may drop this value by only a point or two or less. Upon 1500°F pyrolysis (ambient pressure, inert environment), char yields run around 55-60% giving a total efficiency for liquid phenolic resins of about 40% as a carbon matrix deposition source. During the early 1980s, the most practiced LTV pyrolysis cycle for both ACC and RCC (referred to as 3 day coke pyrolysis) actually covered 4 to 5 complete days (which included cool down) with the parts carefully packed and completely submersed in calcined coke particles. However, a ten hour rapid inert pyrolysis cycle was developed and applied to an ever increasing number of parts. This method, which utilized a constant flow of either Ni_2 or Ar gas throughout the run, was repeatedly applied with success and was on the way to becoming the standard pyrolysis processing cycle. Rapid Inert Pyrolysis saves an enormous amount of time (ie.. part packing, retort preparation, cycle time and cool down), and there were no defects ever observed to be unique to this cycle – it was essentially demonstrated to be equivalent to the longer method. The 3 day process was simply the method employed during the early days and retained ever since.

Micro-structure analysis as well as literature resources have indicated that essentially all matrix materials derived from organic compounds impregnated into the substrate's porosity have one important property in common: the formation of a 'glassy', vitreous inorganic carbon residue upon pyrolysis which resembles the characteristic hexagonal 'benzene-like' structure. This would appear to be true for all polymeric resins as well as pitch residues. Even though this carbonaceous matrix is comprized strictly of inorganic carbon, there is even evidence that 'resonance' electrons are shared throughout the six member ring network.

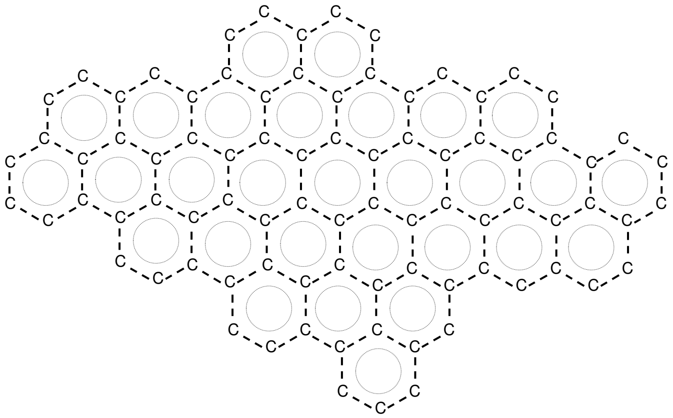

Simplified illustration of carbon/carbon matrix residues when formed from glassy precursors (such as polymers).

While LTV considered a few small scale laboratory experiments related to techniques in chemical vapor infiltration and deposition (CVI, CVD) for carbon/carbon processing, production level densification and coating methods have always been confined to the processes of resin impregnation/pyrolysis and pack mix sedimentation respectively. These approaches significantly set LTV's carbon/carbon technology apart from the rest of the industry which generally practice well established methods of CVI carbon matrix densification followed by SiC CVD surface coating. Morphological differences between matrices formed by pyrolyzed organic deposits (imprenation/pyrolysis) and CVI are notable. Since CVI carbon is deposited or 'grown' at high temperature (1700-2200°F) directly from the carbon source (usually methane), macro-crystals of 'non-glassy' amorphous carbon are typical. Control of the infiltration process is highly dependent on the temperature and pressure employed during the process in order to direct deposition onto the innermost pore surfaces. Compared to the impregantion/pyrolysis approach, matrix densification by CVI is extremely slow since carbon deposition occurs atom by atom and often requires several weeks. In contrast, four cycles of resin impregnation/cure and pyrolysis can be completed in 10 or 12 days when rapid inert pyrolysis is used. Impregnation with coal tar pitch has been explored by researchers and offers benfits as well as drawbacks. A denser layer of aromatic and heavy aliphatic carbon entities can potentially be impregnated into the composite's porosity but these materials will typically liquify and flow out well before they reach carbonization temperature; lower weight components will also tend to volatilize throughout the process. There are environmental and safety risks associated with this method as there are with all the others. A high pressure impregnation-vessel-combination-carbonization-furnace would probably be required to permit an uninterrupted flow of the process from impregnation through pyrolysis under pressure. This would be an interesting option to explore and could reduce densification time enormously.

It is well known that carbon/carbon susbstrates will rapidly decompose when exposed to oxidizing environments above 800-1000°F unless they are rendered oxidation resistant by some method or another. Oxidation protection is a requirement for materials which are to be used for most aerospace applications, particularly the outer panels of vehicles which must re-enter the earth's atmosphere after space missions. It is also well known that the most common method of providing oxidation protection for carbon/carbon substrates is the application of SiC coatings. Without a doubt, the physical nature of the densified substrate (particularly the porosity network) plays an ever increasing role in the degree of success attributed to the coating operation. The uniqueness of LTV's concept of pack mix SiC conversion coating has appeared to offer the most promising approach for producing large exterior space craft panels and is notably different than the predominant industry method utilizing CVD applied SiC surface coatings. For many years, the specific parameters of the coating process and it's exact composition were closely guarded secrets of LTV's long time chief coating technologist (now retired). Precise control of particle packing pressure in a silicon-starved mix as well as the maximum firing temperature and up-ramp play important roles in successful conversion of the outer 15-30 mils (first and second outer plies) of the substrate. The reaction is typically carried out at 2800-2900°F in a mix composition comprized primarily of inert SiC powder so that Si metal (which melts around 2600°F) is compelled to migrate into the substrate to react with both fiber and matrix carbon to form the low porosity, much denser β-SiC ($\rho \sim 3.2$).

Again, the most common method for substrate oxidation protection via surface deposited SiC is practiced by essentially all other carbon/carbon companies in the industry and utilizes the CVD approach. A typical precursor is methyltrichlorosilane (MTS) which decomposes within the CVD furnace and deposits (grows) β-SiC crystals on the substrate in the 1850-2200°F range (α-SiC will form at temperatures below about 1700°F). H_2 gas is also used in the process (perhaps as a spoiler) along possibly with N_2 (as a dilutent). Both gases are generally flowed into reation chamber in the same stream with MTS where the MTS and H_2 crack to begin the deposition process.

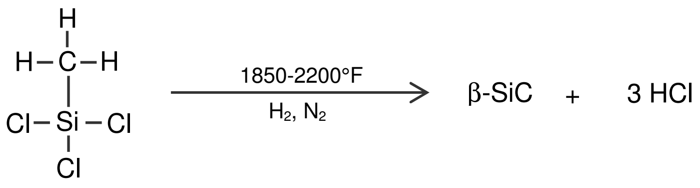

$$CH_3SiCl_3 \xrightarrow[H_2,\ N_2]{1850\text{-}2200°F} \beta\text{-SiC} + 3\ HCl$$

MethylTrichloroSilane

Control of pressure during the run determines the degree to which SiC is deposited either within the internal, accessible pores of the substrate (CVI) or predominantly on the outer periphery of the part (CVD). Process pressures in the 0.1-5 torr range strongly favor infiltration while 10-20 torr and up are often used for CVD. Both temperature and pressure affect the 'residence time' the reacting components are inside the chamber. Obviously, a substrate with little or no surface porosity can only be CVD coated (decreasing the requirement for strong vacuum pumps) but a CVI/CVD treated porous substrate could possibly help alliviate the vast CTE difference that exists between the carbon fiber/carbon matrix substrate and the SiC phase. This method carries with it a high level of environmental concerns and safety hazards; silane side products from the process can be unpredictably explosive and HCl remnant is not always pleasant to face when the furnace is opened. When perfomed appropriately, this approach can fill up and/or close off essentially all of the substrate porosity to produce a non-porous, oxidation protective coating.

Other methods for imparting oxidation resistance to carbon/carbon include the incorporation of inhibitor compounds in the prepreg and densification resins (such as vinyl-o-carborane) and the application of pre-ceramic polymer resin formulations to the densified substrate surface. In the ladder approach, a number of compounds have entered the market in recent years consisting of organic/inorganic polymer systems containing, among other elements, silicon and/or nitrogen in the backbone. These ceramic precursor resins must be crosslinking thermosets in order to properly convert into the corresponding silicon/nitrogen-based ceramic compounds upon firing. A couple of the more notable resin types are the vinyl polyureasilazanes (such as Dupont's Ceraset) and the vinyl polycarbosilanes (Starfire Systems offers some proprietary formulations). Both systems can be diluted with solvents (adjustable viscosity), can contain fillers (such as SiC or Si powders) and crosslink via free radical initiation.

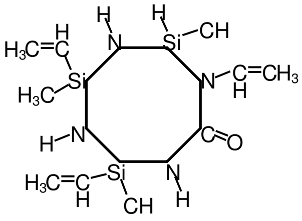

Vinyl ureasilazane monomer (Ceraset)

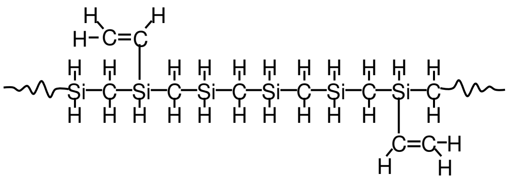

Vinyl polyhydrocarbosilane

While these types of compounds are primarily aimed at the SiC/SiC composites industry, their application as both densification resins and ceramic coating precursors for carbon/carbon has been investigated and employed by other groups in the field. Since I have not fully explored their use specifically in carbon/carbon processing (experience comes from other applications), definitive comments regarding their performance cannot be made here. In general however, firing the carbosilane in an inert atmosphere should produce near stochiometric SiC, while the silazane favors S_3N_4 in N_2 and SiC in Ar and H_2. Depending on the filler make-up, they may tend to form SiC (or Si_3N_4) crystals essentially on the surface of the substrate while custom formulations containing predominantly silicon within the molecule may help assist their potential use as *conversion* coating precusors. Additionally, they have been investigated and employed for crack sealing purposes following the coating process (LTV has traditionally used the liquid glass compound tetraethylorthosilicate TEOS for this process).

An alternative method for the formation of SiC products is the gaseous conversion of carbon (or graphite) directly into β-SiC. This process seems to be scarcely used anywhere within the world. However, it is the primary method utilized by at least one major carbon/graphite company for production of specialized SiC ceramic articles employed by a rapidly growing non-composites related industry. Complete conversion of high density bulk graphite parts up to 1 inch thick is performed on a daily basis. The firing process utilizes simple sand particles and carbon at 3400-3500°F to produce gaseous carbon monoxide and silicon monoxide, the ladder of which attacks the carbon substrate converting it into SiC:

This process has been successfully used on laminated ACC type densified carbon/carbon substrates ¼ inch thick converting the entire carbon fiber/glassy matrix into monolithic β-SiC. Perhaps the most important attribute associated with this method is the formation of extremely high porosity products due to the leaving CO gas. This condition requires an additional process operation (such as SiC CVI/CVD) in order to produce non-porous articles. Obviously, this approach cannot be used to manufacture non-porous, oxidation protective carbon/carbon but could have other applications in the aerospace field requiring SiC ceramics.

Historically, the evidence tends to support the solid state powder-based surface conversion process (unique to LTV/C-CAT) for protection of large aerospace applications of carbon/carbon over most of the other methods. Perhaps the primary benefit of the pack mix approach is related to the nature of the conversion process which produces a distictive *functional gradient* between the two phases by utilizing the composite's peripheral porosity to join the coating to the substrate. A 'fused on' ceramic phase which forms a smooth transition from the SiC medium to the carbonaceous substrate appears to have significant benefits directly related to the CTE and modulus mismatch. Additionally, there is ample evidence that CVD surface coated substrates fail during plasma testing long before conversion coated products regardless of whether the substrate is derived from rayon-based fibers (RCC) or PAN fibers (ACC). However, with respect to the desired levels of CTE compatibility, oxidation protection and substrate mechanical integrity, PAN carbon/carbon is not yet ready to compete on the same level as RCC. . . further development is necessary.

In the late 1980's, one of my collegues from the group left LTV and formed the company called C-CAT of Fort Worth, Texas which has recently been contracted to produced the leading wing edges, nose cap and possibly other panels for the new X-33 space vehicle from standard PAN-based ACC material. These panels are to be incorporated into the vehicle's Thermal Protection System and applied to the most critical hot regions of the ship (the leading edges) similar to the current Space Shuttle. Today, C-CAT's products represent essentially all there is left of LTV's ACC technology. However, this company was established only to *manufacture* carbon/carbon articles utilizing technology concepts relevant to that time period; their charter has never included research, development or technology advancement. Uncoated PAN-based substrates from various companies have found numerous applications over the last decade or so throughout the industry (brakes, clutches, pistons, etc…) but bare substrate is only good to about 900 or 1000°F before oxidative degradation commences. Coated PAN materials, on a small scale, have been employed on a few external space applications (rocket nozzels, body flaps) over the years but these are carefully watched and, quite interestingly, the decision has never been made to replace the Shuttle's RCC with ACC (a fairly simple conversion if the confidence level was high enough). One fact is clear: the current state-of-the-art for PAN-based carbon/carbon is not much different today than is was 20 years ago and it's application to the X-33/RLV vehicle is not really recommended now any more than it would have been back then. While most members of LTV's Advanced Carbon/Carbon Technology (ACT) ACC group were not directly involved with the Space Shuttle's Leading Edge Structural Subsystem (LESS) RCC program (their charter focused on alternative ACC applications for the Air Force and such), it was fairly recognized throughout both groups that PAN-based ACC products which make it through post-coat sealing processes, final inspection and seemingly impressive oxidation testing, carry with them a latent risk. Principal LESS specialists included my supervisor/mentor (the inventor and leading expert), myself (a member of both groups) and one or two other members of the Space Programs division, all of whom have long since moved on or retired. Again, while PAN-based carbon/carbon substrates may look fine after coating (by pack mix, CVD, whatever), a long history of process development and selected applications have repeatedly demonstrated that high modulus ACC components cannot yet measure up to the overall performance, reliability and longivity as the Shuttle's RCC.

Hopefully, these concerns will not be perceived contemptuously, but will remind us all of a few critical problems which have been associated with PAN-based carbon/carbon products since their inception (and likely apply to PAN derived carbon/carbon from almost any source, not just LTV or C-CAT). Even today, NASA is still trying to develop acceptable oxidation protection mechanisms and coating concepts for PAN materials as evidenced in current literature and SBIR contracts. It is the author's professional opinion that modified PAN systems can indeed be formulated that are not only affordable, but can meet or exceed the imminent demands for high tensile strength, higher ILT, coatability, low maintenance and longivity. However, the rayon-based reinforcement approach and pack mix conversion coating methodology are recommended for guidance, and as baseline concepts to be considered seriously since they have provided a performance record that no other technology has yet been able to match. There are some very interesting ideas awaiting to be exploited using these two concepts as starting points.

Newer, specialized surface treatment technologies are now available which could significantly enhance fiber/matrix interactions (including nano particles and nano-sized fibers). If these two phases are not intimately connected throughout the composite network, improvements in interlaminar adhesion can never be fully realized. Alternative multi-dimensional reinforcement approaches should continue to be relentlessly explored (for special applications). For the impregnation/pyrolysis approach, there is no doubt that the correct choice of densification material can maximize the material's cohesiveness as well as its mechanical integrity, while reducing process times.. The matrix precursor must permeate into the core of the composite’s porosity network during each application. Closed pores and voids can and will produce delaminations somewhere during processing or the service life of the product. Porosity must precisely be understood and controlled throughout the production stages. The specific nature of porosity near the periphery of the part (surface porosity) must be well characterized and reproducible in order for the coating process to initiate and progress consistently for each application. A functionally gradient surface conversion approach is an absolure requirement as opposed to the more common surface deposition process. While the pack mix approach is already a proven concept, it is felt that modified CVD/CVI techniques and pre-ceramic polymer coating methods are worth exploring if their goal is aimed at surface conversion rather than surface deposition. There is little doubt that incorporation of substantial phase gradients between the matrix and fiber, and between the coating and the substrate are unbeatable combinations while the CVI densification / CVD coating approach practiced by most carbon-carbon fabricators throughout the world (such as Hitco) have repeatedly demonstrated less-than-desirable properties and often catastrophic results. The apparent solution to the coating/fiber CTE mismatch is to change reinforcement configuration, because the CTE for SiC will always be ~ 4 and *traditional* PAN-based carbon tow will always be around 0 (longitudinally) – but the net modulus of the reinforced substrate can indeed be optimized if the right concepts are applied. With a little willingness to compromise, there are still many unexplored options available . . .

Randy Lee

April 2000