

Hugh McLaughlin, Ph.D., P.E.
151 Hill Road
Groton, MA 01450-1609

tel: 978-448-6066
fax: 978-448-6414
wastemin@net1plus.com
file: IACC14paper.doc

Low Temperature Regeneration of Spent Activated Carbon Theory and Applications

Hugh McLaughlin, Ph.D., P.E.

Text of Paper - see "IACC14figures.ppt" file for Figures

Foreword ⁽³⁾ by Dr. Amos Turk, Emeritus Professor of Chemistry, CCNY

After nearly half a century in Academia studying the applications of activated carbon to real world applications, I have come to appreciate that activated carbon is easier to utilize than to understand. Dr. McLaughlin has undertaken to demystify some of the phenomena observed with activated carbon adsorption and regeneration. While much of what appears in his discussion is not common knowledge, it does represent the current thinking on the inner workings of activated carbon.

Dr. McLaughlin's current paper first develops a coherent framework for understanding activated carbon, and then delves into the low temperature oxidation of activated carbon in air. The conditions in which such oxidation occurs are not well understood and the phenomenon has been unrecognized by the commercial activated carbon industry. McLaughlin has identified the synergism of the slow oxidation of the granular activated carbon and the concurrent and unexpected preferential oxidation of the adsorbed organics. Such a phenomenon provides a novel and attractive method for regenerating spent activated carbon.

After 50 years, it is good to see that there are still new ideas, applications, and regeneration technologies being developed for activated carbon.

Introductory Comment ⁽⁴⁾

The CarbOxLT procedure is the best performing new reactivation/regeneration procedure I have seen for activated carbon in the last 25 years. The performance is comparable to thermal reactivation at 1800F in a steam atmosphere, but uses only a small fraction of the energy/fuel and capital for equipment. ⁽¹⁾

- Dr. Mick Greenbank, Ph.D.

1) Is the CarbOxLT phenomenon new – or was it always there?

While there have been multiple and repeated demonstrations of the CarbOxLT phenomenon to industrially relevant applications ^(1,2,3,4), it is legitimate to question why is this technology just now coming to light. It would seem that something so simple and so fundamental would have long since been identified, explored and commercialized.

The short answer is that most things that will go undiscovered if they are actively being avoided. The CarbOxLT operating environment is adjacent to conditions associated with highly destructive bed fires in activated carbon adsorbers. Thus, there has always been a powerful incentive to stay as far away as possible from the potentially reactive combination of conditions. As such, most studies have been devoted to identifying the onset of oxidation reactions internal to the activated carbon and implementing responses to terminate the “risk”, as opposed to seeking to create and control the core phenomenon for beneficial ends.

The longer answer is that the phenomenon has always been there, and here are some examples of where it is indicated. ANSI/ASTM D 3466-76 (1998), the “Standard Test Method for Ignition Temperature of Granular Activated Carbon” is the measure of when an activated carbon will transition to uncontrolled heat generation in a flowing stream of air. The test consists of gradually heating about one cubic inch of granular activated carbon in a stream of dry hydrocarbon-free air and noting the temperature at which a marked change in the rate of heat rise, associated with the onset of ignition, occurs. The temperature is an indicator of conditions to be avoided for the safe operation of carbon adsorbers. However, one is led to ponder – what is happening as the ignition temperature is approached, yet prior to the onset of the rapid generation of heat?

Mother Nature, the laws of thermodynamics and chemical reactions do not change their minds suddenly at a set temperature. There is an ongoing balance between rate of heat generation and rate of heat removal for any reacting system. The rate of heat generation increases exponentially with temperature and the rate of heat removal increases linearly with the temperature of the heat sink. When the rate of heat removal is greater than the rate of heat generation, the temperature is under the control of the cooler body. If not, then the hotter body increases in temperature at a rate dictated by the difference between the rate of heat generation and removal. Since the hotter body is increasing in heat generation exponentially as the differential temperature increases, as a rule, it doesn't look back until some other constraint is reached (depletion of the fuel source or oxidant, deflagration, detonation, etc.).

As such, the Ignition Temperature represents the transition from the controlled heat removal regime of the ongoing but increasingly exothermic oxidation processes within the activated carbon to the uncontrolled heat generation regime. Before the transition, the process is about as exciting as watching paint dry; after, it turns into a little ball of fire.

The “Ignition Temperature of Granular Activated Carbon” assay demonstrates that either the carbon matrix itself or some other source of oxidizable fuel is present in the activated carbon. The test is explicit that “The test provides a basis for comparing the ignition characteristics of

different carbons, or the change in ignition characteristics of the same carbon after a period of service.” However, as is often the case, ASTM tests are directed to the cautious, not the curious, and generally teach one how to avoid ignitable conditions.

The literature reports the auto-ignition temperature of pure graphite at about 450 degrees Celsius, yet the ignition temperatures of virgin activated carbons are generally in the range of 375 to 425 degrees Celsius, with impregnated carbons exhibiting significantly lower ignition temperatures. Perhaps there is some effect associated with the high internal surface area that characterizes all activated carbons. This feature would result in lower bulk thermal conductivity, as compared to solid graphite, which should shift the heat removal versus generation transition to lower temperatures. With impregnation, the lower ignition temperatures are attributed to (or blamed on) the catalytic effect of the additives.

Unfortunately, the literature is close to devoid of the impact of adsorbed organics on ignition temperature. One Calgon study titled “Assessment of Activated Carbon Stability toward Adsorbed Organics”⁽⁵⁾ does provide some startling measurements of the oxidation threshold of solvent laden activated carbons. That analysis develops a metric called “Solvent reaction (oxidation) initiation temperature (SRIT)” that is utilized to compare the effect of different solvents and loading levels on a spectrum of commercial activated carbons. “SRIT” is defined by a set of experimental characteristics, but represents “an indication of the temperature at which any recognizable oxidation of the adsorbate initiates on a particular adsorbent carbon.”

The Calgon study varied a number of experimental variables, but the most insightful trends are shown in Table 1.

TABLE 1

Organic Solvent	Auto-Ignition T°C	SRIT T°C	(AIT – SRIT)
Toluene	535.6	179.9	355.7
Hexane	225.0	177.9	47.2
Acetone	465.0	124.9	340.2
Methyl Ethyl Ketone	515.6	75.9	439.7
Cyclohexanone	420.0	64.9	355.2

It is apparent from Table 1 that there is no strong correlation between the intrinsic auto-ignition temperature of the pure solvent and the experimental SRIT temperature. Furthermore, the difference between the auto-ignition temperature and the SRIT, representing the depression of the ignition temperature for the solvent upon being adsorbed on activated carbon, is substantial. Furthermore, all tested solvents, when adsorbed on activated carbon, exhibited relatively low temperatures for the onset of oxidation. The dominant effect seems to be some interaction between the solvent and the activated carbon, although the specific mechanism is not apparent.

The Calgon study went on to demonstrate that a cross section of commercial carbons exhibited a 50 degree Celsius variation of SRIT for the same benchmark solvent (MEK). This variation is significant, but secondary to the oxidation enhancement of adsorption on activated carbon

compared to the pure solvent state. Clearly, something about being adsorbed on the activated carbon is dramatically influencing the susceptibility of the adsorbed solvents to oxidation reactions. In addition, the Calgon study demonstrated that the SRIT in air was only 4 degrees higher than the SRIT in pure oxygen. This would imply a minor role for the partial pressure of oxygen in the mechanism of oxidation enhancement and implicate a greater role for some intermediate mechanism, such as chemisorption of oxygen on the activated carbon surface.

There are other examples of conditions that implicate that the CarbOxLT phenomenon is present, although none that materially improve the understanding of the underlying science. The next section will attempt to piece together several observed characteristics to assemble the “state of the art”, which is consistent but disjointed at this point.

2) **Assembling the underlying reactions – reactions of activated carbon alone**

Part of the challenge in understanding the CarbOxLT phenomenon is the intrinsic heterogeneity of activated carbon and the adsorption sites therein, coupled with the spectrum of adsorbates present in most industrial adsorption applications. Furthermore, virtually all organic adsorbates have multiple intermediates that are formed, or at least conceivable, during the incremental oxidation from their original adsorbed form to carbon dioxide and water vapor.

The first piece of the puzzle is the extent of ongoing oxidation of the activated carbon matrix by interaction with molecular oxygen. The phenomenon of low-temperature oxidation of activated carbon in a source of molecular oxygen, typically air, is not new. A body of work performed at the University of Alicante ⁽⁶⁾ in Spain about 25 years ago explored the evolution of previously activated virgin carbon in air at 350 Celsius. As noted in the *Discussion* section of the first paper in Reference 6:

At high percentages burn-off and as a consequence of the weight loss, the adsorptive properties of the active carbons is diminished but they can still be considered good active carbons even when the weight loss is about 50%. This could be important from the industrial point of view, since these carbons could stand temperatures up to 350°C for long periods of time (several days) with no other inconvenience than their reduction in weight.

Additional papers ⁽⁷⁾ further characterized the adsorption characteristics of the activated carbons as air at 350 Celsius was used to modify the internal adsorption sites over a range of up to 70 percent weight loss or “burn-off”. After these four papers, the researchers seemed to drop this line of inquiry and additional work by this group is not reported in the literature.

If the activated carbon is slowly “burning” (for lack of a better term) in air, the question of how long can the “fire” last comes to mind. In laboratory studies, a bed of pure virgin carbon was held at a constant temperature in a steady flow of air and the off gases (predominately carbon dioxide and a minority of carbon monoxide) were measured. Based on the mass balance, the rate of consumption of the activated carbon was calculated. The results are presented in Figure 1, with the results measured for the temperature range of 250 Celsius and above and extrapolated to lower temperatures. The calculated time intervals are the time to consume the entire bed of carbon based on the initial rate of carbon gasification, since the rates would be expected to fall

off as the carbon bed is progressively consumed (first order decay in carbon mass).

As can be seen from Figure 1, the rate of consumption of the bed of activated carbon range from the order of days at the highest temperatures to many years at intermediate temperatures, where the CarbOxLT phenomenon is readily observed, to glacially slow as ambient temperatures are approached. While the ambient predictions are a bold extrapolation, it is comforting to predict that every bed of activated carbon is not just poised to burst into flames.

In theory, the low temperature oxidation of activated carbon shares the same molecular level reactions as the combustion of pulverized coal. Unfortunately, typical coal combustors are operated at higher temperatures, in the presence of more hydrocarbon volatiles, resulting in vapor-phase free radical reactions and significant radiant energy exchange between the reacting surfaces. However, there is some fundamental research in the area of carbon-oxygen reactions that is relevant.

The most significant body of work comes out of the early to mid-1990's research efforts of Professor B.S. Haynes of the Department of Chemical Engineering, The University of Sydney, NSW 2006, Australia. A series of three papers^(8,9,10) explores the "Formation of Metastable Oxide Complexes during the Oxidation of Carbons at Low Temperatures" (title of initial 1990 paper) over the temperature range of 470 K to 720 K (Note: K = Kelvin = °C + 273.15).

Two sources of carbon were used for the Haynes research, a brown coal char and a pure carbon molecular sieve, having micropore surface areas in the range of 600 to 690 square meters per gram. Notably, these surface areas would be significantly less than the typical surface areas for commercial activated carbon, which are generally in the range of 1000 to 1500 square meters per gram. However, the carbonaceous substrates are much closer to activated carbons than typical pulverized coals, both in surface area and in low portion of hydrocarbon volatiles.

The Haynes research concludes that at the lower temperatures studied, complexes of chemisorbed oxygen molecules are formed on the surface of carbon. These complexes may either desorb or react further to form stable surface oxides (C-O bonds) and produce carbon dioxide and carbon monoxide as off gases. At higher temperatures, the predominant reaction is the direct reaction of one oxygen molecule with the graphitic surface to produce one stable surface oxide bond and one carbon monoxide molecule in the vapor phase.

The Haynes research is succinctly summarized by Bews, et.al.⁽¹¹⁾, who note on page 239:

"His studies reveal the importance of the structural heterogeneity of carbonaceous surfaces. Also, O₂ was found to adsorb on carbon in two very distinct ways: in type A adsorption at low temperatures (<500 K) a surface site has one molecule of O₂ attached, whereas in type B chemisorption at higher temperatures (>600 K), O₂ adsorbs dissociatively, giving one atom of O per C site. Type A complexes are not stable at higher temperatures (>600 K): also, a longer exposure to O₂ favors type B chemisorption."

While the Haynes research is very high quality work, it is narrowly focused on elucidating the molecular level interactions between the surface of graphitic carbon and molecular oxygen in the

vapor phase. However, it is significant that the temperature range identified for the presence of the two forms of chemisorbed oxygen (500 K to 600 K, or 227°C to 327°C) corresponds to the range of temperatures where the CarbOxLT phenomenon manifests itself.

3) Assembling the underlying reactions – thermal reactions of adsorbates

While focused studies of individual organic compounds are not found in the literature, it is helpful to review the chemical interactions that one would expect for adsorbed organics under CarbOxLT regeneration conditions, especially in reference to the benchmark behavior of the graphitic matrix of the activated carbon. It is important to factor in the distribution of adsorption energies present in granular activated carbon, which may lead an organic molecule adsorbed in a low energy site to volatilize while the same substance bound in a high energy site may participate in a thermal degradation or oxidation reaction.

As a reference point, consider the activated carbon in an inert atmosphere. The manufacture of activated carbon subjects the graphitic backbone to temperatures exceeding 800 Celsius. In the course of manufacturing the pre-activation calcined intermediate, virtually all the non-carbon atoms contained in organic compounds (i.e. hydrogen, oxygen, sulfur, nitrogen) are thermally removed by volatilization in an inert atmosphere, resulting in the creation of the amorphous graphitic backbone that becomes the activated carbon. As such, it is reasonable to expect that the final activated carbon is essentially non-volatile and will not vaporize to any appreciable extent.

In contrast, organic compounds, upon adsorption, will have a thermodynamic equilibrium between the vapor phase and the adsorbed state within the activated carbon. This equilibrium is highly dependent on the individual organic compound and the adsorption energy levels of the available sites within the activated carbon. However, it is important to recognize that removal of an adsorbed organic under CarbOxLT conditions may occur by direct volatilization of the intact organic compound, in addition to any potential thermal degradation reactions.

Just as the graphitic backbone of the activated carbon is non-volatile, it is also essentially immune to purely thermal transformations and decompositions, since any such reactions have already occurred during the manufacture of the original virgin activated carbon. Note that we are discussing reactions in the absence of oxygen, which precludes the slow oxidation reactions by the graphitic matrix previously discussed

In contrast, in the temperature range of the CarbOxLT phenomenon, it is reasonable and commonplace to see individual organic molecules chemically transformed by dehydration (-H₂O) and elimination (-CO, -CO₂) reactions. Such reactions are especially commonplace in larger organic molecules, particularly molecules containing a high proportion of potential leaving groups, such as carbohydrates.

In summary, under CarbOxLT conditions and in the absence of additional chemical reactants (principally oxygen), the graphitic backbone of the activated carbon can be expected to exhibit negligible volatility and be essentially inert to thermally induced chemical transformations. In contrast, depending on the properties of the organics and binding energy of the adsorption sites,

adsorbed substances can be expected to exhibit both direct intact volatilization and chemical transformations under CarbOxLT conditions.

4) Assembling the underlying reactions – oxygen-adsorbate interactions

Beyond the reactions of the graphitic backbone of the activated carbon with molecular oxygen, there lies the substantially more complex universe of possible reactions of adsorbed organics with the chemisorbed oxygen on the graphite matrix of the activated carbon. In this area, the Calgon “SRIT” study represents the culmination of a short progression of studies focused on predicting conditions that would result in bed fires. Unfortunately, such studies did not attempt to delve into the underlying elementary chemical reactions, but sought merely to identify likely indicator compounds (usually carbon monoxide) that can be used to detect the onset of bed fires in commercial adsorbers before unacceptable temperature gradients are developed.

The first important generalization concerns whether adsorbed hydrocarbons and organics are more or less reactive to chemisorbed oxygen on the graphitic backbone of the activated carbon. By comparing the data of Table 1 and Figure 1, one is led to conclude that the organics are much more reactive, since all the SRIT thresholds are below 200 Celsius and, based on the decade long reaction rate of the activated carbon under similar conditions, one expects that the organics will “burn” orders of magnitude faster.

The final competitive interactions concern whether adsorbed organics will oxidize faster than they will either volatilize or thermally degrade. The case of volatilization (or vaporization) versus oxidation is the conceptually simpler scenario and will be considered first. Due to the large range of adsorption energies in activated carbon, the viable desorption temperatures and equilibrium vapor pressure for an adsorbed organic span a large range for a given organic.

The relationship between the adsorption energy and the equilibrium partial pressure over the adsorbed compound is shown in Figure 2⁽¹²⁾. Figure 3 shows the range of adsorption energies for a representative activated carbon, Calgon CPG, for n-butane, a typical test adsorbate that has been extensively documented in the literature. The butane and propane activity & retentivity thresholds shown on Figure 3 are reference energies for adsorption measured by ASTM tests and are discussed later in the paper.

As shown in Figure 3, this activated carbon has about two thirds of its adsorption pores with adsorption energies above 2.5 kcals per mole of butane adsorbed and 10% of the adsorption sites above 7.5 kcals per mole. Thus, solving the equation in Figure 2, adsorption sites with energies of 2.5 kcals will be in equilibrium with partial pressures of butane of about 230 mmHg. In contrast, 10 percent of the adsorption sites will be in equilibrium with partial pressure of 0.88 mmHg or less.

With such a large range of equilibrium partial pressures, it is likely that a heavily loaded carbon will desorb significant butane when heated and exposed to a vapor stream with a minimal butane partial pressure. It is similarly likely that some portion of the higher energy adsorption sites will remain filled with butane for an extended time. The process described above is similar to steam

stripping, but using air as the stripping medium at a higher temperature. As is experienced in steam stripping, initially a readily desorbable portion of the adsorbed organics is released, followed by a rapid transition to a recalcitrant heel of residual material residing in the higher energy sites.

With respect to recreating future adsorption capacity in the activated carbon, vaporization is a perfectly acceptable mechanism for removing the adsorbed organics. However, vaporization may inhibit any surface oxidation reactions by creating a net outward flux of vapor from the individual carbon particle, thereby inhibiting the diffusion of oxygen into the internal adsorption sites of the activated carbon. Thus, it seems logical that under conditions of pronounced vaporization of organics, the role of any surface oxidation reactions would be reduced.

What is less easy to analyze is the case of equilibrium stripping of partially loaded activated carbon due to under-saturated vapor entering a bed of carbon and exiting nearer saturation. Under those conditions, it is believed that the individual organic molecules move through a packed bed of carbon by chromatographic convection, which involves the iterative sequence of enthalpy-favored adsorption, with localized generation of heat, followed by entropy-favored vaporization. This sequence of “cascading” from site to site down the carbon bed continues until the intact organic exits the carbon bed by convection in the moving vapor phase.

Experimental data indicates that oxidation reactions are promoted under the conditions of chromatographic transport through unsaturated activated carbon. Figure 4⁽¹⁾ shows the data for butane vapors passing through a bed of dry virgin activated carbon for the temperature range of 190 to 230 Celsius. Clearly a dramatic change in oxidation rate is occurring over a relatively narrow and mild temperature range, considering the auto-ignition temperature of n-butane is 405 Celsius.

Figure 5⁽¹⁾ compares the relative reactivity of butane and propane (auto-ignition temperature of 450 Celsius) at 230 Celsius. Propane exhibits a significantly lower exotherm within the carbon bed, compared to butane. Propane would be expected to have less favorable adsorption equilibrium, lower heat of adsorption and lower molar heat of combustion than butane. Any and all three effects could contribute to the trend shown in Figure 5. However, what is known is that neither organic reacts significantly under these conditions in the absence of the activated carbon. Thus, the extent of the “CarbOxLT” phenomenon is variable, but its presence is clearly demonstrated.

In summary, it appears that vaporization and oxidation of organics on the internal surface of the activated carbon are competing parallel processes. Under conditions of pronounced vaporization, the surface oxidation reactions may be suppressed. In the absence of the blanketing effect of organic vapor generation and emigration, the surface oxidation reactions of adsorbed organics are present and may, in fact, be promoted by localized adsorption phenomena and chromatographic transport within the activated carbon mass.

The last piece of the puzzle is the relative reactivity towards thermal degradation versus oxidation for an adsorbed organic. Even the prediction of thermal degradation behavior of adsorbed organics on activated carbon is a sparsely studied research topic, so the comparison of

oxidation versus thermal degradation is essentially pitting the unknown against the unmeasured. However, there are some generalizations about the thermal degradation behavior of organics under the reducing conditions found during high temperature reactivation, thanks to Calgon. Those generalizations are shown in Figure 6⁽¹²⁾.

5) Measuring the extent of temperature dependent reactions

The thermal degradation reactions for the range of temperatures utilized for low temperature regeneration are fairly easy to study. Since thermal degradation reactions often partition the adsorbed molecules into a volatile fraction that is released and a more highly condensed and less volatile residue, the extent of the reaction can be monitored by thermo-gravimetric analysis (TGA). While highly accurate and pricey TGA's are available, the data of interest can be collected with a modified laboratory scale and simple heated sample cell to provide the controlled atmosphere and measure the sample temperature.

Figures 7, 8 and 9 show one such "home-made" TGA. The balance is accurate to 0.1 mg and the temperature of the sample can be measured to 0.1 degrees Celsius. Since the thermal-gravimetric scan is an inherently non-equilibrium process, one critical requirement is to have a highly reproducible temperature ramp, which is provided by an Watlow Anafaze-type programmable ramp & soak controller shown on the right side of Figure 7. The specifics of the TGA will be discussed in greater detail in conjunction with the application of this instrument to the measurement of the adsorption energy distribution of activated carbon.

The intent of the TGA analysis is to provide a repeatable methodology for comparing the weight loss of an individual sample of spent activated carbon under two contrasting atmospheres: inert nitrogen and dry, hydrocarbon-free inlet air. The strategy is to load a sample of spent carbon and stabilize it at a starting temperature, then ascend a reproducible temperature ramp and record the sample temperature and corresponding weight loss versus time.

It is anticipated that, during the nitrogen blanketed runs, the spent carbon will exhibit weight loss due to volatilization of intact organics and thermal degradation of the adsorbed species. This profile of weight loss will be compared with the runs in air, where any oxidation reactions may also occur. However, a TGA scan is specific to the individual spent carbon and the adsorbed organics thereon; moderation needs to be applied before expanding any observed trends to a more global conclusion.

With a TGA scan alone, it is not possible to distinguish between organics volatilized and those generated by thermal cracking of larger adsorbed molecules. Performing off gas analyses could make this distinction, but that would greatly expand the analytical effort and associated expense. Some insight is gained by merely being in the lab, since most thermal degradation products are small, partially oxidized molecules such as ketones and aldehydes – which are easily detected by the human nose. However, principal objective is to isolate the incremental impact of the presence of oxygen versus the background thermal processes involving only the spent carbon.

Figure 10 shows the TGA scans for three representative industrial spent carbons in nitrogen, referred to as CFS3, CII and NA. CFS3 is a vapor-phase pelletized carbon from an aromatic

hydrocarbon removal process, while CII and NA are liquid-phase carbons from sweetener purification applications.

Figure 10 shows the setpoint temperature profile and the corresponding sample temperatures. The sample temperature is offset and lags the setpoint temperature at all times. This allows the sample temperature to respond to any internal exotherms generated within the spent carbon sample.

One feature of the TGA scans in Figure 10 merits discussion. While the CII and NA scans are the same basic shape, the CFS3 scan is different. Based on the origin of the spent carbons, it is hypothesized that the CFS3 scan represents a carbon that is losing weight principally via volatilization and CII & NA show characteristics indicative of thermal degradation reactions (dehydration and rearrangement reactions to yield small volatile organics as leaving groups). These observations are compatible with the broader evaluation of the TGA behavior of organics, in oxygen-free nitrogen, detailed in Reference 13. Reference 13 also develops a methodology for predicting the partitioning between volatilization and thermal degradation with char formation

Figure 11 shows the TGA scans for same three spent carbons in nitrogen and also in air. For both the CII and NA spent carbons, the TGA in air scan is dramatically different, with evidence of an increased rate of weight loss in the middle of the temperature scan, corresponding to samples temperatures in the range of 200 to 250 Celsius. In contrast, the CFS3 sample produced essentially the same TGA scan in both air and nitrogen atmospheres. The effect is shown in Figure 12; where the incremental air effect is compared to the baseline nitrogen scan. As before, CII and NA show significant differences due to air and CFS3 is essentially unchanged, as indicated by the green “delta wt % - CFS3” line on the abscissa.

The behavior of the three spent carbon samples is consistent with the hypothesis that carbon samples that are actively volatilizing intact adsorbed organics, as in the case of CFS3, inhibit oxidation reactions. In contrast, the carbons undergoing thermal degradation reactions demonstrate strong evidence of oxidation reactions in addition to the presumed background degradation reactions under the same conditions. With the data available, it cannot be determined whether there is a synergistic effect between thermal degradation and oxidation reactions with the current TGA experimental results.

Figure 13 shows the data of the two thermally degrading and oxidizing carbons (CII and NA) with additional data on the temperature scans. It is noteworthy that the most reactive carbon, CII in air, exhibits a noticeable temperature exotherm at the time of the most rapid weight loss. This effect provides additional credence to the presence of oxidation reactions, since the thermal degradation reactions are significantly less exothermic compared to oxidation reactions. Close inspection of the data in Figure 13 reveals a similar but less pronounced exotherm for the NA carbon at 40 minutes, corresponding to about 270 Celsius.

In summary, the TGA methodology provides a simple and relatively quick screening of the responses of a spent activated carbon sample to low temperature regeneration conditions. The responses are basically thermal effects, such as volatilization and thermal decomposition of adsorbed organics and oxidative effects, attributed to the “CarbOxLT” phenomenon. Based on

the overview of results presented, it is apparent that response of individual spent carbons can vary significantly due to variations in the amount and properties of the adsorbed organics.

6) Exploring the energy distribution of adsorption sites in activated carbon

Granular activated carbon is a unique material that combines intrinsic properties of naturally occurring starting materials and aggressive activation conditions to create a micro-porous graphitic lattice that possesses exceptional adsorption properties. The adsorption properties are largely due to the polarizability of the localized graphitic structures, which results in “London dispersion forces” that facilitate pure physical adsorption within the internal surface area of the activated carbon.

Figure 14⁽¹²⁾ depicts the “Tug of War” that represents the equilibrium between the adsorbed state in the activated carbon particle and the disassociated state in either a liquid solution or gas phase. Figure 15⁽¹²⁾ depicts the thermodynamic criterion for adsorption from a liquid solution. Figure 2, discussed previously, depicts the same criterion for adsorption from the gas phase.

Both equations assume that the adsorbed state is modeled by a concentrated liquid phase being formed within the pores of the activated carbon. As such, when the temperature and concentration in the bulk phase reach saturation (C_{sat} in the liquid phase or P_{sat} in the vapor phase), the concentrated liquid phase is in equilibrium with the concentration in the bulk phase and the additional potential energy necessary for adsorption is defined as zero (natural log of 1 equals zero). For adsorption at liquid phase concentrations lower than C_{sat} and gas phase partial pressures less than P_{sat} , the necessary additional potential energy required to achieve adsorption can be calculated.

As a conceptual starting point, envision the internal structure of activated carbon filled with an adsorbed vapor. Figure 16⁽¹²⁾ depicts a small internal region of activated carbon for the case of butane adsorption from a stream of pure butane at room temperature. Under such conditions, much of the micropore structure is flooded with liquid-like adsorbed butane, but some regions have much higher adsorption forces attracting an individual adsorbed butane molecule. This broad range of adsorption sites that form a distribution of adsorption energies characterizes all adsorption applications using activated carbon.

The challenge is to measure the distribution of adsorption energies within a given sample of activated carbon under conditions that provide meaningful insight to the industrial application one is striving to understand. In general, activated carbon is typically characterized by adsorption from the gas phase, since activated carbon reaches equilibrium with gas phase concentrations very quickly, facilitating rapid execution of the analytical procedure.

For the purposes of understanding this new low-temperature regeneration technology, several analytical procedures have been utilized that measure adsorption properties of the activated carbon for known vapors under controlled conditions. These methods are used to compare one carbon against another, with virgin activated carbon representing one logical reference point. By using these methods and testing multiple carbons subjected to various regeneration and

reactivation conditions, insight into the effects of reactivation and regeneration on spent activated carbon can be accumulated.

Using the equation contained in Figure 2, at a fixed temperature, a wide range of adsorption potentials can be simulated by varying the pressure, and thereby the vapor phase concentration, of the adsorbable compound. Figure 17⁽¹²⁾ shows the range of pressures required for butane adsorption at -0.5 Celsius (the boiling point of butane at one atmosphere) to simulate the entire range of adsorption forces present in a typical granular activated carbon. As can be seen, almost ten orders of magnitude in partial pressure are necessary to map the entire adsorption potential distribution.

When one does map the entire adsorption potential distribution for a given sample of activated carbon, a distribution such as shown previously in Figure 3 is generated. Figure 3 shows the relative amount of adsorption capacity at a given adsorption force for Calgon's CPG activated carbon product. Given that the heat of vaporization of butane is 3.5 kcal/mole, the heat of adsorption for the highest energy sites in activated carbon is greater than four times the heat of vaporization for butane. As such, adsorption in activated carbon can be viewed as "facilitated condensation", where the London dispersion forces shift the equilibrium from the bulk vapor phase to a condensed liquid phase in the pores of the carbon.

Because of the difficulty in measuring the entire adsorption force distribution, many analytical tests measure one or two intermediate adsorption energy levels and quantifying the volume of adsorption capacity above those energy thresholds. One such test is the Butane working capacity test, ASTM D5228. This method takes a dried sample of carbon at 25 Celsius and measures the weight gain upon being equilibrated in pure butane at one atmosphere, as depicted in Figure 16. The weight gain is termed "Butane Activity". The sample is then purged with 1000 bed volumes of dry air over 40 minutes and the residual weight gain is measured and termed "Butane Retentivity". The difference is termed the "Butane working capacity".

De facto, the Butane Activity measures total adsorption capacity above a relatively low adsorption energy threshold and Butane Retentivity measures the capacity of those adsorption sites above an energy level that represents roughly one third to one half of virgin activated carbon capacity. While the total adsorption capacity measured by Butane Activity is worth knowing and correlates well with other measures of total pore volume (BET surface area, iodine number, CCl_4 number), a significant portion of this capacity is too low in adsorption energy to be of any practical value in an actual industrial adsorption application.

Since butane working capacity methodology can be performed with any challenge gas, other gases will measure different adsorption energy thresholds. Challenge gas is the historical term for a test vapor that the carbon is "challenged" to adsorb. Less easily adsorbed gases provide measures of correspondingly higher adsorption energy thresholds. Two convenient alternatives to butane are propane and R134a (1,1,1,2 tetrafluoroethane, the HFC replacement for automobiles). These alternatives tend to quantify activated carbon capacities at equilibrium with the pure vapor (the "Activity" measure) with adsorption energies that are sufficient to be of utility in industrial applications.

Figure 3 shows the estimated adsorption energy thresholds for butane activity and retentivity, as well as the corresponding thresholds for the same assay performed with propane. Figure 18⁽³⁾ shows the adsorption capacity distributions for a number of popular industrial carbons, using esoteric historical units advocated by Calgon.

While it may be tempting to conclude that all carbons look the same, based on Figure 18, it should be noted that small differences on a semi-log plots often represent significant performance differences in the real world. In addition, the curve labeled “React”, signifying a high-temperature reactivated carbon, is discernibly below the other curves in the lower adsorption potential range. This indicates lower total adsorption capacity, as measured at the propane activity threshold, which is commonly observed in actual adsorption applications with reactivated carbons.

In summary, the adsorption capacity of either virgin or regenerated activated carbon can be assayed to determine the available adsorption capacity. The measurement can consist of either one or two single point measurements, which quantify the total adsorption capacity above a specific adsorption energy threshold, or the entire “characteristic curve”, which maps the entire distribution of adsorption capacity as a function of adsorption potential.

7) Tracking the CarbOxLT regeneration using Propane Working Capacity

While the background rate of oxidation for virgin activated carbon has been measured as slow to very slow, as discussed in conjunction with Figure 1, the rate of oxidation of spent activated carbon, under similar conditions, is significantly more rapid. A sample of spent virgin activated carbon, heavily loaded with sugar refining color bodies, was regenerated over a total of about 24 hours at progressively higher temperatures. The results of the propane working capacity assay, performed at each temperature plateau, are shown in Figure 19.

As can be seen in Figure 19, at each progressively higher regeneration temperature, incremental recovery of adsorption capacity is measured (Note: the sum of Propane Retentivity and Working Capacity equals the Propane Activity).

Figure 20 shows a comparison of high-temperature reactivation and low-temperature regeneration, performed on the same batch of spent carbon used for Figure 19. Figure 20 includes the properties of dried spent carbon and spent carbon that has been calcined (heated to 850 Celsius in an dry nitrogen atmosphere, which reduces any adsorbed compounds to char). Virgin carbon values for propane activity, propane retentivity and density are used to normalize the data shown in Figure 20, hence all virgin carbon properties are shown as 100%.

As can be seen in Figure 20, the dried spent carbon shows a significant increase in density and loss of propane capacity compared to the virgin carbon. Calcining the spent carbon recovers a significant portion of the adsorption capacity, but also converts the initial adsorbed material into non-volatile char, which then accumulates over repeated adsorption-calcining cycles and incrementally decreases the available adsorption capacity, as observed in Figure 20.

Calcining followed by steaming simulates the high-temperature reactivation process, since the steam removes the char by the water gas reaction. This allows the density of the reactivated carbon to be returned to near virgin carbon density, yet the propane capacities are still significantly decreased from the original virgin carbon levels.

Finally, the results for the low-temperature regeneration are shown on the far right of Figure 20. The data shown in Figure 20 is for a different regeneration of the same spent carbon as Figure 19 and the second regeneration results were slightly better than before. However, compared to the starting virgin activated carbon, the propane capacity decrease for the low-temperature regeneration is roughly one half that experienced with high-temperature reactivation.

While it cannot be demonstrated using the data presented, low-temperature regeneration conditions should result in significantly less attack on the graphitic backbone of the activated carbon than high-temperature reactivation. The milder conditions should reduce the generation of fines during carbon transport, which is largely attributed to the mechanical weakening of the carbon particles due to repeated reactivations. As such, for a circulating pool of activated carbon, the make up rate with virgin activated carbon is anticipated to be less for low-temperature regeneration than would be required for a similar system using high-temperature reactivation.

In summary, based on measurements of density, it has been demonstrated that the low-temperature regeneration technology is capable of controlling the bulk density of the regenerated carbon pool at near virgin carbon levels. In addition, based on measurements of propane activity and retentivity, the low-temperature regeneration consistently recovers a greater portion of the adsorption capacity as compared to the same spent carbons being subjected to high-temperature reactivation.

However, referring to Figures 3 and 18, it is clear that the entire “characteristic curve” of a given activated carbon cannot be unequivocally inferred by just two data points measured by propane activity and retentivity. The next section will look at the entire “characteristic curve” to gain further insight to the low temperature regeneration process.

8) Mapping the upper range of adsorption energies with GRPD

As discussed earlier in conjunction with Figure 17, roughly ten orders of magnitude in pressure variation are required to cover the entire range of adsorption energies in activated carbon for butane at a constant temperature (which also fixes the saturation pressure, P_{sat}). However, if one varies the temperature, then a much broader range of adsorption energies can be covered over a more accessible range of experimental conditions.

Figure 21 ⁽¹²⁾ shows the variation of adsorption potentials for butane at one atmosphere from – 0.5 Celsius to 550 Celsius. This temperature range covers the same range of adsorption energies as the ten orders of magnitude previously depicted in Figure 17. Furthermore, by choosing a less easily adsorbable compound, the temperature range can be further narrowed.

Figure 22 shows the relationship of temperature and adsorption potentials for butane and R134a

(1,1,1,2 tetrafluoroethane, the HFC replacement for automobiles). R134a is a relatively difficult compound to adsorb, especially at elevated temperatures. However, it forms a relatively dense condensed phase upon adsorption, which facilitates the analytical challenge of measuring the amount of R134a adsorbed by weighing. Thus, pure R134a vapor over the temperature range of 150 to 300 Celsius is in equilibrium with the upper half of the adsorption potentials in activated carbon.

By comparison of Figures 18 and 22, it can be inferred that the adsorption potential for adsorption of R134a at 150 Celsius is essentially the same as the energy threshold for the propane retentivity test. Thus, the R134a assay allows the energy levels above the propane retentivity to be quantified and augments adsorption capacity characterization that can be collected from the butane and propane working capacity tests.

Measuring the uptake of a “challenge gas” over a range of temperatures is termed “Gravimetric Rapid Pore Determination”, which is an analytical technique evolved via a series of improvements over the prior art discussed in References 14 & 15. The method is very powerful, in that it measures the adsorption energy associated with each incremental amount of adsorption capacity, but it is not really a standardized analytical technique. It is basically a research tool that is best applied when comparing one carbon to another. It is in this semi-quantitative role that this analytical procedure will be used here.

Gravimetric Rapid Pore Determination (GRPD) instruments are basically custom thermogravimetric analyzers, but there is no accepted standard instrument or procedure. Basically, a sample of carbon is subjected to a controlled environment of temperature and gas composition, and the relative weight change relative to a standard condition is measured. However, since only the upper portion of the adsorption capacity is of interest, the incremental weight changes are typically at the limit of analytical detection.

The apparatus utilized in these studies is essentially the same apparatus used for the Thermogravimetric analysis discussed previously and shown in Figures 7 to 9. When performing a GRPD scan, the temperature scan is modified in an effort to measure the adsorption capacity of the activated carbon under equilibrium conditions.

Figure 7 shows the general layout of the custom instrument, consisting of a laboratory scale with resolution of 0.1 mg and the ability to weight below the scale, a gas supply consisting of a pressurized source of dry nitrogen and the challenge gas (R134a), which is regulated through a rotameter, and a custom temperature controlled housing that is heat traced with laboratory tapes and controlled by a digital temperature controller.

Figure 8 shows the “boat” that holds the activated carbon sample within the uniform temperature housing. The boat has a sintered metal frit for a bottom, is open on top and holds approximately 1 gram of activated carbon (roughly two cubic centimeters). A thermocouple is located just above the boat within the housing and provides the temperature measurement for the temperature controller and for data acquisition. Figure 9 shows the readout of the scale and the temperature indicator connected to the thermocouple within the housing.

In operation, the apparatus is loaded with a carbon sample and heated to above 300 Celsius with dry nitrogen flowing vertically up the temperature-controlled housing (see Figure 7). The sample weight is allowed to stabilize and the scale is zeroed at this setting. The atmosphere is switched to the challenge gas, the heat input is reduced and the temperature of the housing is allowed to gradually descend to 150 Celsius, which is the temperature range of interest for these studies. Below 150 Celsius, the heat input is increased and the apparatus returns to above 300 Celsius over approximately the same time period, roughly 30 minutes per scan.

The ongoing weight readings of the scale and the concurrent temperature readings are recorded. The data is then tabulated by noting the weight measurement every 10 degrees on both the descending and ascending traces. The average of the two weights corresponding to each temperature is used as the weight gain of challenge gas at that temperature. When the apparatus is unloaded, the boat is emptied of activated carbon and the scale reading noted for the boat once it has cooled to room temperature. This weight equals the dry weight of the activated carbon, before the challenge gas was started, as established when the scale was zeroed at the start of the run.

The analytical capability described in above is potentially fraught with pitfalls with respect to the quality of the data generated. Any such apparatus should be subjected to repeated test runs to establish baseline stability and reproducibility. It should also be calibrated with known standards of virgin activated carbon. When these standardizing runs were performed for the apparatus shown in Figure 7 to 9, a significant baseline drift was observed for a sample of virgin carbon in dry nitrogen over the temperature range of interest. It was determined that over the course of the temperature ramp, the changing temperature of the rising air current from the housing was shifting the zero measurement on the scale. An air wash, consisting of a small computer cooling fan and ductwork to direct the ambient air past the entrance to the balance, stabilized the baseline. A supplemental fan, shown in the upper right of Figure 7, further isolated the thermal effects of the housing from the scale.

After debug of the analytical apparatus, one can move onto the surprises contained in the samples of interest. One such unexpected result is shown in Figure 23, where the entire adsorption curve for identical samples of virgin activated carbon seems to shift with the range of temperatures measured. The effect is genuine, and not due to experimental error – the underlying phenomenon is the “annealing” of the internal surfaces of activated carbon.

Annealing, when used to describe this phenomenon, is the process by which the non-carbon atoms present within the activated carbon sample are removed by heating the sample in an inert atmosphere. Activated carbon has a well-documented tendency to chemisorb oxygen and also tightly adsorbs water molecules in the higher energy pores. When the carbon is heated in dry nitrogen, as is done at the start of the GPRD scan, the carbon sample loses moisture and desorbs surface bound oxygen in the form of carbon dioxide and carbon monoxide. The removal of water vapor is dependent on both time and temperature, assuming a bone dry desiccating vapor phase, whereas the release of surface oxides is predominately a function of temperature, as discussed in References 16 and 17.

Removing the surface oxygen groups has two pronounced effects on the adsorption performance

of activated carbon. Removing the oxygen groups increases the volume of individual pores and also increases the dispersive (non-polar) energy of adsorption of the enlarged pore, as discussed in Reference 18. The net effect is to increase both the total adsorption volume and average energy of adsorption, resulting in a better performing adsorbent. Unfortunately, the improved performance is basically an analytical anomaly and creates difficulties comparing analytical test conducted at room temperature, such as the butane and propane assays, with test performed at elevated temperatures, such as GRPD.

Transformations being initiated by the elevated temperature conditions utilized during the GRPD assay becomes much more pronounced when analyzing samples that have been partially regenerated by the low-temperature technology. Since the CarbOxLT technology subjects the surface of the activated carbon to atmospheric levels of oxygen at elevated temperatures, the chemisorption of oxygen is facilitated. Furthermore, many of the intermediate regeneration samples have never been subjected to temperatures as high as 300 Celsius, so the potential for additional thermal decomposition reactions by adsorbed compounds exists when the sample is subjected to the GRPD temperature range.

All these simultaneous phenomena make the definitive analysis of a GRPD assay of an activated carbon sample a difficult task at best. However, the core objective is to scan the range of adsorption energies to see if high-temperature reactivation samples exhibit substantially different trends than samples regenerated by the CarbOxLT technology. To the extent the analytical conditions for each sample are identical, the differences observed can safely be attributed to differences in the samples and the associated regeneration methods.

Figure 24 shows the GRPD scans for the samples discussed in conjunction with Figures 19. The same data is plotted in Figure 25 as a semi-log plot. Because of the linear relationship between temperature and adsorption potential for R134a, as shown in Figure 22, and the relatively constant density of the condensed R134a phase over the temperature range of measurement, Figure 25 is essentially an equivalent characterization as Figure 3, differing only in the units used to scale each axis.

Figure 25 clearly shows the uniform progression of the regeneration process, impacting on all sizes of pores and fully developing even very high-energy adsorption sites. Figure 26 shows the trends of Figure 25 plus duplicate scans for virgin activated carbon GAC1240 (the Norit product). It is notable that the virgin carbon measures slightly lower adsorption capacities than the CarbOxLT sample regenerated to 325 Celsius, although the propane retentivity measurements of Figure 19 would predict the opposite trend. This is attributed to the annealing of the 325 Celsius sample, which can increase the total adsorption volume and energies, as previously discussed.

Figure 27 shows the GRPD scans of the samples discussed in conjunction with Figure 20. The trends are consistent with the propane retentivity measurements with the exception of the CarbOxLT regenerated sample, which has shifted to essentially virgin carbon properties. This pattern is reasonable, since the calcined and high-temperature reactivated samples have both been subjected to temperatures much higher than the GRPD conditions, implying that any effect of annealing would be evidenced in both the propane retentivity assay and the GRPD scan.

In summary, the GRPD studies indicate that low-temperature regeneration impacts the same broad range of adsorption sites as the high-temperature reactivation, including the highest energy sites that would be associated with removal of trace levels of hard-to-adsorb chemicals, such as taste and odor compounds. However, the GRPD analytical technique is a great deal of work and not a straightforward assay to execute or interpret. As such, the facile and representative guidance afforded by the propane working capacity assay is recommended for process optimization measurements.

9) Comparing the energy requirements of Spent Carbon Reuse Options

When regenerating spent activated carbon, there are material losses and the industry practice is to provide additional virgin carbon to “make up” the carbon lost during regeneration. One can easily appreciate how the make up rate per cycle of additional activated carbon impacts on the cost of operating a carbon regeneration system, since the cost of purchasing the make up virgin carbon quantifies the cost. However, make up carbon costs rarely dominate the overall cost of operation of a carbon system, except in the case of “use once and discard” processes, such as with landfilled granular and powdered activated carbon. There are other costs that are less obvious, which range from direct fuel costs for the regeneration equipment to depreciation of the capital investment of the regeneration facility.

Calgon published an article years ago that does a good job of delineating the economics of operating a high-temperature reactivation system⁽¹⁹⁾. Unfortunately, since the article dates from 1978, the utility and capital costs are not accurate for today’s industrial climate. However, the overall structure of the economic analysis is still sound and comprehensive.

Relative to the high-temperature reactivation technology, the low-temperature regeneration approach has two significant advantages that impact on the economics of operation. The first advantage relates to direct operating costs, specifically the fuel requirements. As discussed in the Calgon paper, the fuel consumption for high-temperature reactivation is several thousand Btu per pound of carbon processed. If an afterburner is required to destroy the volatiles generated during the charring of the adsorbates, the fuel requirement may well double. The high fuel consumption of reactivation and high operating costs are a direct result of operating at high temperatures and using the “water-gas” reaction, which is endothermic – even at 850 Celsius.

Figure 28 shows the enthalpy increase or the energy input to heat the spent activated carbon up to the reaction temperatures for low-temperature regeneration and high-temperature reactivation. In Figure 28, the enthalpies of all the components of the spent carbon (retained liquid water, the granular activated carbon and gases such as nitrogen and oxygen) are set to zero at room temperature (77 Fahrenheit). The steam curve in Figure 28 is elevated relative to the water curve by the heat of vaporization required to boil the water, since somewhere in the industrial complex, cold make up water was supplied to the boilers that generate the steam for the reactivation process. Thus, Figure 28 shows the heat requirement to reach the reaction temperatures. To calculate the heat duty, one starts with the incoming spent carbon and raises the temperature of the individual constituents to the required reaction temperature, including energy requirements to

vaporize any liquid water present.

As a rule of thumb, drained carbon retains about one pound of liquid water for each pound of dry granular carbon. The enthalpy requirements of adsorbed compounds can be estimated by partitioning the compounds into liquid water, based on the amount of available H₂O, and treating the rest of the adsorbed material as carbon or GAC. Thus, for example, sucrose (C₁₂H₂₂O₁₁) would be treated as eleven water molecules and twelve carbon atoms, or 58 weight percent water and 42 weight percent GAC. If supplemental air or steam is supplied to the reaction, the enthalpy required to raise the starting material (ambient air or make up water to the boiler) to the reaction conditions can be estimated using Figure 28.

While heat balances are a tedious accounting exercise, it is possible to calculate the adiabatic energy requirement to heat the reactants to the reaction temperatures. However, one can appreciate by looking at Figure 28 that the energy input to raise one pound of moist spent carbon to CarbOxLT temperatures is clearly less than to reach the high-temperature reactivation zone. Furthermore, high-temperature reactivation requires one-third to one-half pound of steam per pound of dry reactivated carbon to react with the char (and carbon backbone), which is a significant heat input at the high temperatures already required. To make matters worse, the water-gas shift reactions are endothermic to the tune of about 5000 Btu's per pound of graphitic carbon reacted away, so incremental additional heat is required as the water-gas shift reaction proceeds.

In addition to lower heat requirements to reach reaction conditions, low-temperature regeneration is exothermic and can use the heat generated by the oxidation of the adsorbed organics to supply the bulk of the necessary heat to maintain the reacting mass at the desired regeneration temperature. Probably the lowest heating value organic compounds that adsorb on activated carbon are carbohydrates, as represented by sucrose in the sugar refining industry. Since sucrose has a heat of combustion of roughly 7000 Btu's per pound and since most adsorbed compounds are very similar to sucrose in their ratio of carbon to non-carbon atoms, the heat of combustion of most adsorbed organics in sugar decolorization is similar to that of sucrose.

When one does the heat balances, one finds that wet spent activated carbon, containing one pound of retained water and 0.2 pounds of adsorbed color bodies per pound of dry carbon, reacts under CarbOxLT conditions as an "autogenous" reaction – where the heat of combustion matches the heat requirement to heat incoming spent carbon to the reaction temperature, including vaporizing the retained water. Above 20 weight percent color bodies, a level that includes virtually all spent decolorizing carbons, the fuel value of the color bodies being destroyed supplies the energy requirement for the low-temperature regeneration process. As such, the direct fuel costs for the low-temperature regeneration process are estimated at less than one-quarter of the fuel costs for high-temperature reactivation.

For hydrocarbon adsorbents, the heat of combustion is significantly higher than for carbohydrates such as sucrose. As such, the adsorption loading required for autogenous reaction conditions is less. Table 2 shows the chemical reactions and estimated heat of combustion for typical hydrocarbons, with separate values provided for aliphatic and aromatic hydrocarbons. One finds that drained liquid phase carbon (one pound of water per pound of dry carbon) that has

8% by weight (dry weight of carbon as the weight basis) adsorbed hydrocarbons will regenerate autogenously under CarbOxLT conditions. For vapor phase carbons, which generally still have 10 percent adsorbed water, the autogenous fuel requirement is as little as 2 percent adsorbed hydrocarbons.

10) Estimating the Capital Investment requirements of Carbon Reuse Options

Estimating direct energy costs is much easier than estimating capital costs to install either regeneration technology. Capital costs are more dependent on individual installation features, due to variations in labor costs and local requirements for ancillary processes, such as air pollution abatement equipment. Furthermore, since there have not yet been any full scale implementations of the low-temperature regeneration process, it is hard to develop a credible estimate to compare the comparable capital costs for multiple hearths used by the high-temperature reactivation process. However, it is possible to put the two technologies in perspective.

The cost of the high-temperature reactivation process is sparingly documented in the literature and it is not cheap. Reference 20, dating from 1989, estimates \$3,650,000 for a ten million pound per year regeneration capacity. Reference 19, by Calgon and dating from 1978, estimates \$2,700,000 for a 30,000 pound per day reactivation rate (almost 11 million pounds for a 365 day operating year), of which \$770,000 is purchased equipment and \$1.93 million dollars is “Installation” costs. Reference 21 provided an estimate for \$980,000 for just the core multiple hearth unit of 10 million pounds per year of reactivated carbon. Thus, it looks like the cost of the multiple hearth capability does not come cheap, which can be confirmed by asking almost anyone who owns one.

While the exact cost of an equivalent regeneration capacity using the low-temperature technology has yet to be demonstrated, it is expected to be significantly less than the high-temperature process for one simple reason: operating temperatures and the corresponding acceptable materials of construction. High-temperature reactivation requires refractory lined equipment and exotic alloys due to the high operating temperatures, 800 to 1000 Celsius. Low-temperature regeneration operates at temperatures of less than 400 Celsius, which allows for carbon and stainless steels without linings to be used for equipment fabrication. Thus, due to dramatically different operating temperatures, not only are the initial construction capital costs significantly reduced, but the operational flexibility and reliability is greatly improved over refractory-lined designs.

There is a concern that the low-temperature process is five to ten times slower than high-temperature reactivation, so one might expect the low-temperature equipment to be larger to provide for sufficient reaction time. However, it turns out that the low-temperature equipment is actually predicted to be of similar size, due to the different “volumetric efficiency” of the two technologies. Multiple hearth furnaces are only about 10 percent full of carbon, due to the activated carbon being contained on a series of shelves, called hearths, in layers that average only a few inches deep. In contrast, the low-temperature process utilizes free-flowing “dense beds” of carbon, meaning that the reactor volume is more completely filled with carbon. These

beds are mixed to control the temperature of the reacting mass and avoid hot spots.

There are additional unanswered concerns relating to the construction and operation of a full-scale low-temperature regeneration facility, which fall in the general category of “fear of the not-yet-known”. However, as Dr. Mick Greenbank, Ph.D. of Calgon so kindly summed it up ⁽⁴⁾:

The only unanswered question is with reliability and maintenance, but that cannot be proven till the first systems are installed. However I would be surprised if the CarbOxLT regeneration systems were not more reliable than existing high-temperature reactivation technology.

In summary, it appears that the adopters of the low-temperature regeneration technology may save a significant amount of money when building the facility and will save a significant amount of money while operating the facility. At that point, it comes down to the individual organization’s engineering capabilities, tolerance for risk, and propensity for saving money.

11) What does a Commercial CarbOxLT Reactor look like?

No one knows, since one has not been built yet. However, there are pilot-scale reactors in operation and they have shown the feasibility of regenerating spent carbon on a commercial scale.

Figure 29 shows the outside of the heat shroud of a laboratory scale reactor, which has a volume of 4.4 liters and regenerates approximately 1 kilogram of carbon at a time. Figure 30 shows the heat shroud removed and inside rotating drum that holds the activated carbon. During regeneration, the drum rotates slowly and the carbon cascades down the angle of repose, as visible through the glass dome, traveling from upper left to lower right in Figure 30. The laboratory reactor has the air introduced at the far end of the reactor and the vapor stream travels back toward the glass dome and flushes out the annulus around the entrance of the air injection tube into the reactor

Figure 31 shows a large pilot-scale reactor, capable of regenerating approximately 25 pounds of spent activated carbon per batch. The basic features are similar to the laboratory reactor and the batch conditions and cycles times are fundamentally the same. The pilot scale reactor has a foil-enshrouded catalytic converter, shown in the upper center part of Figure 31. The catalytic converter oxidizes volatile organics that vaporize intact and small organic compounds resulting from thermal decomposition of the adsorbed compounds. In addition, it converts any carbon monoxide to carbon dioxide before discharge to the atmosphere. One of the primary purposes of the pilot scale reactor is to allow measurement of emissions from the regeneration of a representative sample of spent activated carbon from an operating industrial process. For spent carbons that do not have appreciable halogenated adsorbed compounds, the vapor emissions are not anticipated to require any further treatment prior to discharge to the atmosphere.

Figure 32 shows the basic details of a commercial-scale reactor. The reactor can be operated in either a batch mode or with continuous feed and discharge. It is proposed that the basic design of Figure 32 can be scaled up to daily capacities as high as 10 tons (or more) of regenerated carbon per day.

12) Summarizing the Low-temperature Regeneration Process

The major differences between low-temperature regeneration, CarbOxLT (LTR) and High-temperature reactivation (HTR) are summarized in Figure 33. While both processes allow the recycling of spent activated carbon adsorbent, they are really two entirely different approaches to the same problem.

The core technical differences between the current multiple hearth reactivation process and the low temperature regeneration technology comes down to time, temperature, and reaction conditions. The high temperature reaction conditions first convert all the adsorbed organics into residual char in the carbon pores. Then, the high temperature process uses steam to react with the char (and the activated carbon itself). This reaction occurs rapidly above 850 Celsius and is endothermic, which means it adsorbs heat (versus generates heat). The high temperature reaction is controlled by adding heat to maintain a temperature setpoint and metering the steam supply.

The low temperature reaction occurs at between 150 and 350 Celsius, reacts oxygen (from air) directly with the adsorbed organics, and generates heat (exothermic). The reaction is also relatively slow. In addition, the temperature is controlled to prevent the development of hot spots, which might otherwise combust the activated carbon directly. The low-temperature reaction is controlled primarily by removing heat, either by heat exchange or by injecting water, to control a temperature setpoint and, to a lesser extent, by metering the air input.

Another important difference is the dominant vapor atmosphere present within each process. The low-temperature process has excess oxygen present at all times and the predominant off-gases are carbon dioxide and water vapor, with lesser amounts of carbon monoxide. In contrast, high-temperature reactivation uses the water-gas shift reaction and generates an atmosphere containing principally carbon monoxide and hydrogen. These compounds require subsequent destruction, typically with an afterburner. In addition, due to the presence of hydrogen, there are safety concerns and the risks of further reactions with reactive vapors present in the furnace off-gas, such as those that form hydrogen cyanide.

13) Comparing High Temperature Reactivation with CarbOxLT Regeneration

Since many industrial applications for activated carbon regenerate and recycle carbon adsorbents to “extinction,” the cumulative effect of multiple loading and regeneration cycles is important. A laboratory study was performed to simulate the evolution of “pool” carbon properties. The details of this laboratory study are discussed in Reference 3. In summary, the laboratory study utilized virgin activated carbon (Norit GAC1240, another typical sugar decolorizing carbon) and simulated the adsorption step, under laboratory conditions, by equilibrating the carbon with a

controlled ratio of industrial molasses.

During the equilibration step, the color bodies from the molasses adsorbed into the activated carbon, which was then “sweetened-off” with water rinses and regenerated. This sequence was repeated several times, without the addition of any make-up virgin carbon, in order to establish the trend over multiple adsorption-regeneration cycles.

Figure 34 shows the propane activity trend of multiple adsorption-regeneration cycles for virgin GAC, low-temperature regenerated carbon, high-temperature reactivated carbon and activated carbon that was regenerated by calcining. The regeneration of spent carbon by calcining alone reduces the adsorbed organics to char, which does recover some adsorption capacity. However, the char accumulates in the regenerated activated carbon and progressively consumes adsorption capacity. As such, calcining alone does represent the reasonable “worst-case” regeneration, just as restoring the spent carbon to virgin activated carbon properties represents a reasonable “best-case”.

As shown in Figure 34, the calcining alone of spent carbon results in a progressive loss of adsorption capacity, as measured by the uptake of propane defined by the “Propane Activity” assay discussed previously. High-temperature reactivation is basically just calcining followed by the “water gas” reaction. As can be seen, the steep loss of propane activity previously seen with calcining alone is avoided, although the capacity of the high-temperature reactivated carbon still declines a significant amount with each adsorption-reativation cycle.

Figure 34 also shows the trend for low-temperature regeneration. For each adsorption-regeneration cycle, the low-temperature method was used to return the bulk density of the regenerated carbon to near-virgin carbon levels. Low-temperature regeneration results in a noticeably slower rate of adsorption capacity loss, as measured by propane activity, than high-temperature reactivation. However, the loss of capacity compared to virgin activated carbon is still discernable. As such, low-temperature regeneration is not a “perfect regeneration”, but rather a distinguishable improvement over the current industrial practice of high-temperature reactivation.

Figure 35 shows the same comparison for the three regeneration techniques using propane retentivity as the measure of the condition of the regenerated carbons. The overall trends are similar to the trends demonstrated by the propane activity assay, but the differences between the various regeneration methods is considerably less pronounced.

In closing, it is acknowledged that the actual specifics of a low-temperature regeneration operation will not be known until one or more full-scale units are constructed and operated for a sufficient period of time. However, it is clear that the low-temperature regeneration process enjoys several pivotal technical advantages over the high-temperature reactivation technology that currently represents the industrial “status quo”.

References:

- 1: “CarbOxLT – The State of the Art – Autumn 2002”, Hugh McLaughlin, Ph.D., P.E., presented at the 10th International Activated Carbon Conference, Pittsburgh, PA, September 26, 2002. An electronic version of this paper may be obtained by contacting the author.
- 2: “Method for destruction of organic compounds by co-oxidation with activated carbon”, Hugh Stanley McLaughlin, inventor, U.S. Patent Application 20030206848, filed April 29, 2003. See also Patent Cooperation Treaty WO 03/106330, Application number PCT/US03/13516.
- 3: “A New Technology for Regenerating Sugar Decolorizing Activated Carbon”, Hugh McLaughlin, Ph.D., P.E., presented at SPRI 2004 - Sugar Processing Research Conference, April 4-7, 2004, Atlanta, Georgia, USA. See proceedings of this conference or contact the author for information on obtaining a reprint of this paper.
- 4: “Exploring the Subtleties of Low-temperature Regeneration of Sugar Decolorization Activated Carbon”, Hugh S. McLaughlin, Ph.D. of Hugh McLaughlin, P.E., Groton, MA, USA & Stephen J. Clarke, Ph.D. of Florida Crystals Corporation, South Bay, FL, USA, presented at the Annual SIT Conference, May 16-19, 2004, Vancouver, B.C., Canada. See proceedings of this conference or contact the author for information on obtaining a reprint of this paper.
- 5: “Assessment of Activated Carbon Stability toward Adsorbed Organics”, Edmond C. Akubuiro and Norman J. Wagner, I&EC Research, Vol. 31, pp 339-346 (1992).
- 6: “The Controlled Reaction of Active Carbons with Air at 350°C – I: Reactivity and Changes in Surface Area” P. Gonzalez-Vilchez, A. Linares-Solano, J. de D. Lopez-Gonzalez and F. Rodriguez-Reinoso, Carbon Vol. 17, No. 6, pp. 441-446 (1979) and “The Controlled Reaction of Active Carbons with Air at 350°C – II: Evolution of Microporosity”, F. Rodriguez-Reinoso, A. Linares-Solano and J.M. Martin-Martinez, Carbon Vol. 22, No. 2, pp. 123-130 (1984).
- 7: “Adsorption of hydrocarbons on air-reacted activated carbons. I. Adsorption isotherms”, J.M. Martin-Martinez, A. Linares-Solano, F. Rodriguez-Reinoso & J. D. Lopez-Gonzalez, Adsorption Science & Technology, Vol. 1, pp. 195-204 (1984) and “Adsorption of hydrocarbons on air-reacted activated carbons. II. High and low pressure hysteresis”, A. Linares-Solano, F. Rodriguez-Reinoso, J.M. Martin-Martinez & J. de D. Lopez-Gonzalez, Adsorption Science & Technology, Vol. 1, pp. 317-327 (1984).
- 8: “Formation of Metastable Oxide Complexes during the Oxidation of Carbons at Low Temperatures”, A.E Lear, T.C. Brown, and B.S. Haynes, Twenty-Third Symposium (International) on Combustion, The Combustion Institute, pp. 1191-1197 (1990).
- 9: “Oxygen Chemisorption on Carbon”, T.C. Brown, A.E Lear, and B.S. Haynes, Twenty-Fourth Symposium (International) on Combustion, The Combustion Institute, pp. 1199-11206 (1992).
- 10: “Surface Heterogeneity in the Formation and Decomposition of Carbon Surface Oxides”, M.C. Ma and B.S. Haynes, Twenty-Sixth Symposium (International) on Combustion, The

- Combustion Institute, pp. 3119-3125 (1996).
- 11: "The Order, Arrhenius Parameters, and Mechanism of the Reaction between Gaseous Oxygen and Solid Carbon", I.M. Bews, A.N. Hayhurst, S.M. Richardson, and S.G. Taylor, *Combustion and Flame*, Vol 124, pp. 213-245 (2001).
 - 12: Slide from a course titled "Selecting the Best Activated Carbon for a Specific Application", presented September 2002 by Mick Greenbank, Ph.D., of Calgon Carbon Corporation. The course was presented by Dr. Greenbank while in the employment of PACS, a specialty activated carbon service firm located in the greater Pittsburgh area (www.pacslabs.com). Many such slides have text references to "CCC", which stands for Calgon Carbon Corporation.
 - 13: "Study of Thermal Regeneration of Spent Activated Carbons: Thermogravimetric Measurement of Various Single Component Organics loaded on Activated Carbons", M. Suzuki, D.M. Mistic, O. Koyama and K. Kawazoe, *Chemical Engineering Science*, Vol. 33, pp. 271-279 (1978).
 - 14: "Flow Method for Determination of Desorption Isotherms and Pore Size Distributions", B.P. Semonian and M. Manes, *Analytical Chemistry*, Vol. 49, No. 7, pp. 991-994 (1977).
 - 15: "New Test Methods for the Activated Carbon Industry", M. Greenbank, H. Nowicki, H. Yutes and B. Sherman, *Water Conditioning & Purification*, pp. 92-96, February 2003 issue (available online at <http://www.wcp.net/PDF/NewTestMeth02-03.pdf>).
 - 16: "Characterization of Oxygen-containing Surface Complexes created on a Microporous Carbon by Air and Nitric Acid Treatment", Y. Otake and R.G. Jenkins, *Carbon*, Vol. 31, No. 1, pp. 109-121 (1993).
 - 17: "The effect of the gradual thermal decomposition of surface oxygen species on the chemical and catalytic properties of oxidized activated carbon", G.S. Szymanski, Z.Karpinski, S. Biniak and A. Swiatkowski, *Carbon*, Vol. 40, No. 14, pp. 2627-2639 (2002).
 - 18: "Application of Inverse Gas Chromatography at Infinite Dilution to study the effects of Oxidation of Activated Carbons", J. Jagiello, T.J. Bandosz and J.A. Schwarz, *Carbon*, Vol. 30, No. 1, pp. 63-69 (1992).
 - 19: "Selecting a thermal regeneration system for activated carbon", R.H. Zanitsch and R.T. Lynch, *Chemical Engineering*, January 2, 1978.
 - 20: "Cost Estimates for GAC Treatment Systems", J.Q. Adams and R.M. Clark, *Journal AWWA*, pp. 35-42 (January , 1989).
 - 21: Private communication. Stephen McDonough, Vice President, Sales and Marketing, Hankin Environmental Systems Inc. (www.hankinenv.com).