



How Activated Carbon Works to Purify Air and Water

How does Activated Carbon Work to Purify Contaminants for Liquid and Vapor Phase Applications?

The first step is to define the performance limiting factors in the application. For this application, most of the adsorber is used for MTBE adsorption in the ppb concentration range. Adsorption of BTEX, TBA, or humic acids or other total organic carbon (TOC) components are removed by the front end of the column. In theory, any multi-component adsorption application can be split up into multiple objectives and more than one carbon selected. For example, an optimal MTBE carbon on the effluent end of the column and a different optimal TOC removal carbon on the influent. In practice, few customers optimize their carbon selection to a level that justifies application of multiple carbons. If only one optimized carbon is applied, then it should be optimized for the species utilizing the greatest amount of carbon. In this application that species is the ppb level MTBE.

Each Starting Material Results in One Calcined Structure

The definition of activated carbon is a crude graphite plate structure, with some localized order on a molecular level. The adsorbing pores are defined by the gap, angle, and spacing between the graphite plates. The raw materials used for activated carbon production, such as coal, coconut shell, and wood do not have any graphitic structures, and therefore no significant volume of pores for adsorption. But each starting material has all the chemical building blocks, which create the size and shape of the graphite plates that are formed once the temperature is above 1,450°F. Above 1,450°F the most stable form of carbon is graphite and most carbon bonding will rearrange to form graphite plates.

Calcination is defined as the heat treatment of the raw material above 1,450°F and in an inert atmosphere (to assure no further activation processing occurs). During calcination most all heteroatoms are removed and all carbon atoms are either lost as volatile matter or are converted to crude graphite and become part of the calcined carbon structure. A specific starting material can generate only one calcined structure with a single density, which is the maximum possible with that starting material. In other words, the calcined structure has the maximum number of carbon possible with the carbonation chemistry of the starting material. The calcined structure is also a fixed rigid structure due to the interbonding between the plates. The chemical building blocks of the starting material also dictate the nature of this interbonding between the plates, and thus the gap, angle, spacing between the plates, which dictates the carbon pores structure.

Concept of a Family of Activated Carbon for a Starting Material


There is only one family of activated carbon pore structures that can be made from a given starting material. The calcined structure gives a single structure and maximum apparent density (AD) for a specific raw material. Unfortunately this calcined structure has so little density, and has so little porosity, that it has limited commercial value. Additional porosity, or adsorption sites, need to be created by removing carbon atoms from this fixed rigid structure, leaving voids, which are available for adsorption. The only mass in an activated carbon particle or granule is carbon. The carbon atoms are removed from inside the particle and the shape and size of the granules remain constant. Therefore as carbon atoms are removed the AD of the granules decreases in direct proportion. The AD of a carbon product specifies the number of carbon atoms remaining and the AD of the product divided by the AD of the calcined product (AD_{prod}/AD_{calc}) determines the fraction carbon atoms remaining and is a measure of porosity created. Also one minus this ratio ($1 - AD_{prod}/AD_{calc}$) is a measure of the porosity created during the activation process. A raw material can generate a single family of products, from the single calcined structure, through a variety of larger porosity carbons. The individual members of the family are distinguished by the AD specification. The amount of

porosity developed, is specified by the carbon tetrachloride test. Under the conditions of this test, the molecule is easily adsorbed and every site within the structure that is available for adsorption is filled. Its value is proportional to the total number of adsorption sites, or total adsorbing pore volume.

Development of Porosity during the Activation Process

Activation is a process that converts the calcined structure into a family of activated carbon products at lower AD values. The activation process is simply the ordered removal of carbon atoms from inside the granule to create porosity. The removal is orderly and predictable because the most chemically reactive carbon atoms in the calcined structure are removed first. Therefore at a given time and temperature, a specific number of carbon atoms will be removed from the most reactive sites, and the AD will be decreased to a specified value, and a certain member of the raw material product family will be produced, defined by its AD. Low AD products have the highest porosity, or highest activity.

There are two reactions that are preferred to make a homogeneous product using a thermal activation process:



Both of these reactions are endothermic, causing instantaneous cooling near the site where the carbon atom was removed. This decreases the probability that these cooler carbon atoms will be attacked next. These endothermic reactions eliminate the 'fuse' effect, or uncontrolled burning, which occurs with exothermic reactions such as:



Therefore the activation equipment must have good atmosphere control to allow good gas-solid contact between the CO_2 and H_2O and the graphite plates, but prevent any oxygen (air) from seeing carbon structure. If the activation is not homogeneous, then there is a potential for the product to be a mixture of high-density regions mixed with low-density regions, and the product AD would be a weighted average. This is referred to as heterogeneity and generally is to be avoided, as a mixture of AD products with an average AD rarely outperforms the homogeneous product of the same AD.

Basis of the Adsorption Force Present

Adsorption forces are due to intermolecular interactions with carbon atoms in the structure. These intermolecular attractions are called London forces and analogous to gravitational forces but on a molecular scale (they are additive temperature independent and nonspecific). The adsorption force for a given void, or adsorption site within the structure, is the sum of all the interactions with all the nearby carbon atoms. The adsorption properties carbon can be fully characterized by simply measuring the distribution of numbers of sites with different adsorption force (adsorption potential or energy is actually easier to measure but is closely related).

Figure 1 shows the statistical distribution of number of sites at the various adsorption forces or potentials. This distribution can be fit to a mathematical form with three empirical coefficients listed on the graph. These characteristics define the adsorption properties of the activated carbon. Figure 2 shows how these distributions and these three parameters change with AD as a raw material is activated. This change in distribution of sites follows a simple principle. The more carbon atoms in the structure, the more intermolecular interactions, the greater the overall adsorption force, averaged over the entire structure. As mentioned before the AD defines the number of remaining carbon atoms and therefore the average adsorption force over the entire structure. High AD carbons are more strongly adsorbing, averaged over all the entire structure, than low AD carbons.

Therefore it is always a compromise in activated carbon production, to create additional porosity by decreasing the AD removing more carbon atoms, or to leave enough carbon atoms behind to provide the

intermolecular interactions, required to generate strong adsorption forces. The pores don't adsorb anything; it is the interactions with the carbon structure. Yet if the entire granule is carbon structure, the adsorption forces will be very strong, but there will be a shortage of porosity to place the adsorbed molecules.

The best AD, or best performing member of a product family, is determined by the requirements of the application. Low concentrations of smaller molecules with higher solubilities (or vapor pressures) require more adsorption force to remove and only the higher AD products have significant adsorption sites with the forces required. For large molecules with high concentrations and low solubilities (or vapor pressures), any AD carbon will work, so low AD carbons are used because their additional porosity can hold more molecules. Performance prediction programs can model any gas or liquid phase application and define the exact adsorption force requirements to aid in the selection of the optimal product AD. The effect of the extent of activation (the AD of the product family member) on the adsorption isotherm is shown in Figure 3. The isotherm shape would have been the same for liquid-phase adsorption of MTBE, as for gas-phase adsorption of butane; the x-axis would simply be ppm concentration instead of partial pressure. At high concentrations (near saturation or solubility limits) the low AD product is preferred and at low concentrations the high AD outperforms.

Six Types of Applications

The amount of adsorption force required can be calculated exactly, from the physical properties of the adsorbing molecule and the concentration in solution. The capacity on the carbon is the total number of adsorption sites that have adequate adsorption force. The sites with inadequate adsorption force remain empty and are not a factor.

In the real world, applications involve complicated mixtures of a wide variety of constituents. But from doing a lot of performance prediction calculations it was found that the world of activated carbons was actually simpler than originally believed. All applications could be grouped into six types based on their carbon requirements. Figure 4 shows the relative performance curves for three of these application types for the family of a specific coal-based product line. The optimization curves have broad optima, and six of these curves (application types) will cover the entire range of performance with near optimal performance. A long theoretical discussion can prove that almost all, physical adsorption applications can be classified into one of these six types based on the adsorption force requirements and the portion of the carbon structure filled with adsorbate. The MTBE adsorption applications from aqueous solutions between 5 and 500 ppb happen to be Type IV.

Selecting the Family Member with Optimal Performance

Figure 4 shows the range of products available from the specific coal raw material and the relative performance for Types I, III, and IV applications. Each point on this graph is a different activated carbon structure, a different product family member for this coal, defined by the AD on the x-axis.

The best carbon for Type I applications is the lowest AD possible. In commercial applications, a product with an AD of 0.35 g/cc is used because the lower AD products are too expensive and too fragile to be practical. This is because manufacturing costs, yields, and other physical properties, like particle strength, also vary with AD. Lower density means longer activation times, or higher temperatures are required to remove more carbon atoms, all which result in greater expense. The AD decreases with the removal of carbon atoms also causes the structure to become more fragile (analogous to wood and termites).

Type III applications the optimal member of the coal based product family is an AD of 0.6 g/cc. For Type IV application the optimal product is 0.7g/cc, but the commercial product sold with this raw material is 0.6 g/cc. The reason is kinetics. At very high AD values there is so much carbon structure and so little porosity that the rate of adsorption is slowed to the point of affecting performance. One solution is to use a smaller particle size, such as 30x70 or 20x50. The particle diameter is roughly proportional to the average diffusion path length molecules have to travel to adsorb. This relates to the residence time (empty bed contact time) the stream has to have, in contact with the carbon, to allow adsorption. In other words, the slower the

adsorption the longer the residence time required, and the larger the adsorber size (and amount of carbon on line) will be for a given flow rate. The Type IV curve in Figure 4 is broad. The difference in performance between 0.66 and 0.7 g/cc is minimal, but the difference in adsorption rates is significant. Therefore a 0.66 g/cc AD coal-based product is the optimal family member for Type IV applications.

Selecting the Starting Material with the Member with Optimal Performance

Figure 5 shows the optimization curves for a variety of raw materials: a specific West Virginia coal, coconut shells from the Philippines, cherry pits from Michigan, peach pits from California, pecan shells from Texas, and palm kernels from Nigeria. Each point on this graph is a different activated carbon structure, and each curve is the entire family of products that can be produced from the individual raw materials. To select the best raw material the optimal carbon from each raw material family of products is compared. In Figure 5 the best product is the palm kernel product at 0.6 g/cc, but unfortunately this product is not currently being manufactured on a commercial scale. Peach pits looks favorable but the calcined density is about 0.4 g/cc, so to reach the higher AD values to get the higher performance would require costly carbon deposition processing.

The next option is the 0.55 AD coconut product, which demonstrates a 35% performance advantage over anything that could be made with the coal-based starting material. Due to adsorption rate considerations and empty bed contact time requirements, the preferred commercial coconut product is on the low AD portion of this curve, closer to an AD of 0.50 g/cc. (A smaller granule diameter such as 30x70 mesh size which also helps minimize the empty bed contact time).

Actual Performance Predictions

Table 1 shows the results of performance for multi-component calculations for MTBE and TBA and BTEX components for the most popular liquid phase coal-based carbon at 0.48 g/cc in commercial applications and the selected coconut-based product for this applications at 0.50 g/cc. The table reports the use rates in lbs. of carbon required to purify 1,000 gallons of water. The use rates are cumulative, meaning the use rate is the carbon required to remove that component and all other more strongly adsorbed components in the stream. It is obvious that MTBE is the critical component, and its requirements dictate the carbon selection from 0.01 to 1 ppm (10 to 1000 ppb) MTBE concentrations.

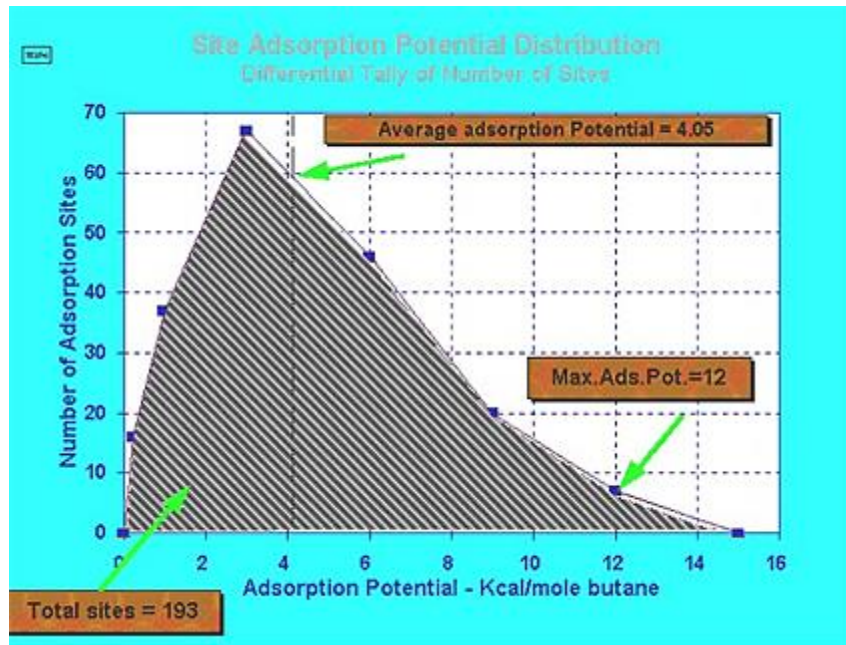


Figure 1. The three key characteristics defining distribution of adsorption sites by adsorption force

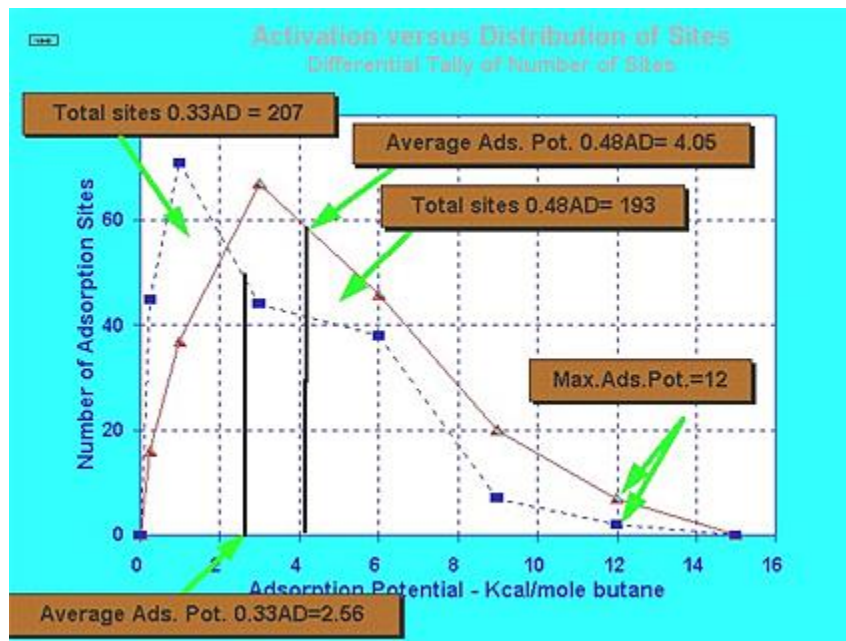


Figure 2. Change in distribution of adsorption sites with activation to lower apparent density (AD)

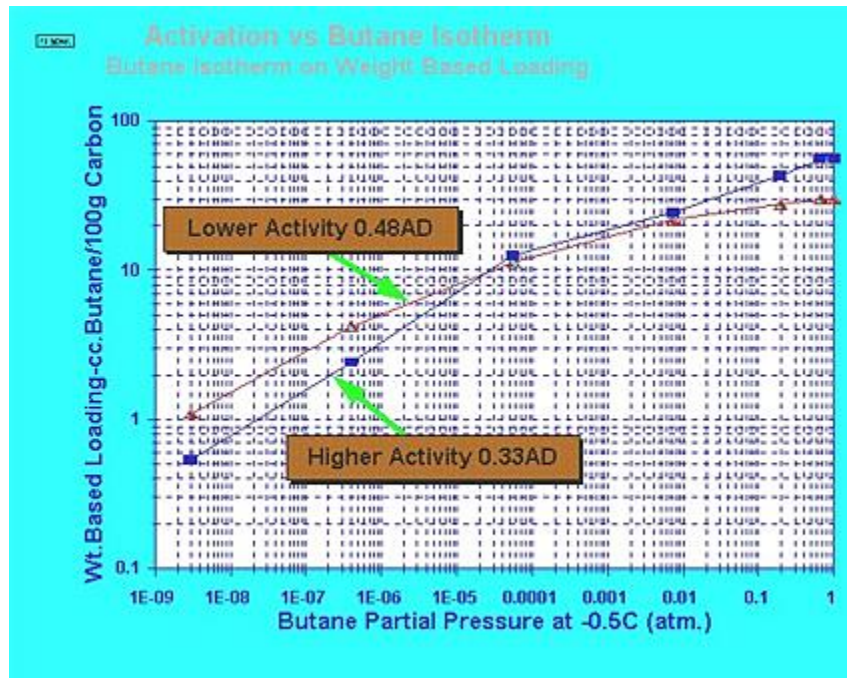


Figure 3. How the change in distribution of adsorption sites with AD affects adsorption isotherms

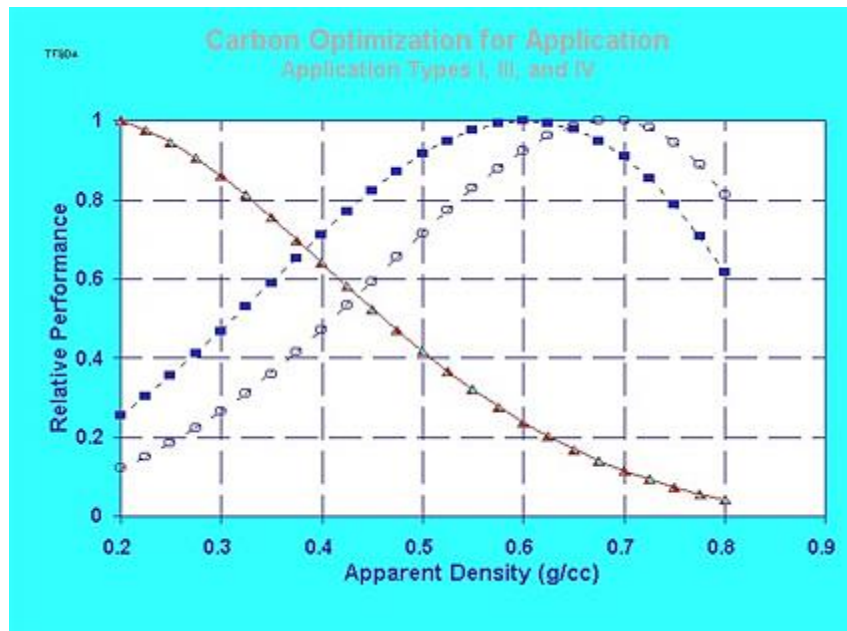


Figure 4. Comparison of relative performance of three of the six types of adsorption applications

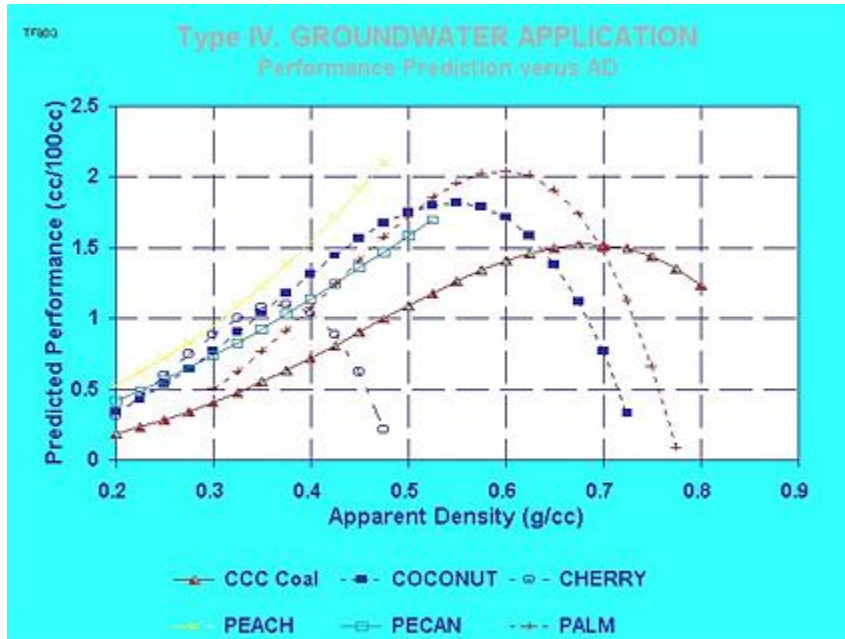


Figure 5. Comparison of performance of product families from different raw material - Type IV applications