

The many faces of carbon in electrochemistry: general discussion

Robert Dryfe, George Zheng Chen, Domnik Bayer, Martin Jonsson-Niedziolka, Richard McCreery, Julie Macpherson, Héctor Abruña, Yury Gogotsi, Chi-Chang Hu, François Béguin, Laurence Hardwick, John Foord, Matteo Duca, Elizabeth Howe, Katherine Holt, Manuel Alvarez-Guerra, Patrick R. Unwin, Thomas Rabbow, Milo Shaffer, Jet-Sing Lee, Elzbieta Frackowiak, Fulian Qiu, Boris Dyatkin, Philip A. Ash, Keith Stevenson, Dirk Guldi and Philippe Bergonzo

DOI: 10.1039/C4FD90025A

George Zheng Chen opened the discussion of the Introductory Lecture by Richard McCreery: In your list of electrode capacitance for different carbons, the values differ significantly. Would you please comment on the causes of such differences, and also explain how you have measured the capacitance of the respective carbon materials? Particularly, what electrolyte have you used in the measurements?

Richard McCreery replied: We used either 0.1 M HClO₄ or 1 M KCl in water for capacitance measurements, and used cyclic voltammetry, where the charging current is proportional to the capacitance and the scan rate. HOPG has an unusually low capacitance.

George Zheng Chen opened the discussion of the paper by Héctor Abruña: It is well known that conducting polymers can be used for various electrochromic purposes, such as displays and smart windows. What are the improvements of your graphene–PEDOT hybrid in comparison with previous work? Also, if your graphene–PEDOT hybrid is to be used for smart windows, what is the light transmittance?

Héctor Abruña responded: You are certainly correct in that there is well-established literature on the use of conducting polymers for electrochromic applications. Part of the intent with the graphene–PEDOT hybrid that we have prepared was to move away from ITO-type electrodes, which have been traditionally employed in these applications. Graphene offers the advantages of being conductive and highly transmissive. (Single-layer graphene is about 97–98% transmissive.) In addition, it can make a conformal coating. In terms of light transmission, since the graphene itself has very low absorbance, the transmittance of the composite will be dominated

by the thickness of the PEDOT film. One can readily make a calibration curve of transmittance *vs.* polymer thickness so as to adjust and optimize the PEDOT film thickness to the particular application/need.

Richard McCreery asked: Have the oxygen containing functional groups on graphene been characterized, for example by XPS? You invoke them to explain the dependence of electronic properties on pH, but has anyone established the identity of the oxygen containing functionalities?

Héctor Abruña answered: We have not done XPS characterization of the oxygen functionalities on the graphene surface. My sense is that such measurements are likely to be difficult because the signal levels are likely to be low and broad. We invoke their presence as a (reasonable) way to rationalize the pH dependence.

Julie Macpherson queried: There are different types of oxygenated functional groups that will form on the graphene surface. For pH FET sensing is there a preferred oxygen functionality for designing an optimal pH sensor?

Héctor Abruña responded: Although we do not have any direct information concerning the various oxygen functionalities that may be present on the graphene surface, my sense is that carboxylic acid groups are likely to be the best for pH sensing. I base this on the fact that such groups (as opposed to hydroxyl or aldehyde groups) are more likely to respond rapidly and reversibly to changes in pH.

Katherine Holt asked: What is the required density/coverage of surface functional groups to allow for pH sensing in FET devices? Has anyone calculated or determined such a value?

Héctor Abruña replied: While to my knowledge nobody has calculated or determined this, one can estimate the value. The saturation coverage of quinones on HOPG is of the order of 1.5×10^{-10} mol cm⁻².¹ Given the voltammetric profiles and S/N considerations, we would estimate (clearly speculative on our part) that a surface coverage of 1–5% should be sufficient to provide pH dependent behavior.

1 M. T. McDermott and R. L. McCreery, *Langmuir*, 1994, **10**, 4307.

Katherine Holt remarked: With regards to the pH sensitivity of functional groups on graphene, our paper¹ discusses the pH dependent behaviour of graphene nanoflakes with carboxylic acid edge groups. The oxygen functionalities of these nanoflakes are well defined, which allows us to ensure the pH dependent behaviour can be attributed to the carboxylic acid functionalities. We haven't used the graphene nanoflakes for pH sensing but it is possible they would be a good candidate material for this application.

1 M. M. Lounasvuori *et al.*, Electrochemical characterisation of graphene nanoflakes with functionalised edges, *Faraday Discuss.*, 2014, DOI: 10.1039/c4fd00034j.

Héctor Abruña added: Graphene-type materials, such as graphene nanoflakes with a high density of oxygen functionalities, would appear to be good candidates

for such an application. My concern would be how they would be supported on a conducting substrate to which one could make contact.

Milo Shaffer communicated to Richard McCreery and Héctor Abruña: The issue of oxide structures on carbon has been studied in various ways. Recently, I heard an excellent presentation by Neil Wilson (Warwick) describing a detailed study of atomic oxygen adsorption on graphene, associated with atomic resolution microscopy and *in situ* ARPES measurements. He showed the complexity of the resultant structures and the fact that a couple of O atom% were sufficient to destroy the band structure of the graphene seen in ARPES. It is interesting to compare this result with the estimate given by Richard McCreery for the maximum 1% functionalisation before the whole layer is affected. In then considering Hector Abruña's devices, and his estimate that >1% functionalisation is required for a successful pH sensor, are the key properties of graphene lost too soon? Or is the associated change in properties relevant to the device operation?

Héctor Abruña communicated in response: I was not aware of the study by Neil Wilson, thanks for bringing it up. The issue of 'oxygenated groups' on carbon surfaces has been (and continues to be) the subject of numerous studies. The 1% estimate that I gave was based more on S/N considerations for measurement. In the statement, "the fact that a couple of O atoms were sufficient to destroy the band structure of the graphene", what would this correspond to in terms of surface coverage? As carbon surfaces go, graphene is likely to be the most affected in terms of the consequences of having oxygenated species. A detailed study in which the surface coverage, and defect density, are controlled would be of great utility. However, this is likely to be very difficult.

Richard McCreery communicated in reply: I agree on most points. Everyone should be careful to distinguish between electronic disturbances and actual chemically distinct sites. Outer sphere systems should be sensitive to electronic effects, which may extend well beyond the site of the defect itself. An electrocatalytic process, however, may depend on specific functional groups present only at the defect. An example is aquated $\text{Fe}^{+3/+2}$, which is catalyzed by surface C=O groups.

Patrick Unwin asked: How does the response in Fig. 11 of the paper for bare graphene compare to that for HOPG? Additionally, what happens to the response on the graphene electrode in subsequent cycles? I ask this as the electrochemical oxidation of NADH can produce products which foul the electrode. Is this seen on graphene?

Héctor Abruña answered: While we have not done an exhaustive study of changes in the response for NADH oxidation during consecutive scans, there are some observations that we have made. For the oxidation of NADH at bare graphene, we did observe a decay in the magnitude of the response with consecutive scans. For graphene surfaces modified with 3,4-DHB, we did not observe dramatic changes. While somewhat speculative on our part, we believe that this could be due to depletion of NADH within the diffusion layer, as well as potential blocking/inhibiting effects for the formation of non-reactive products.

Philip A. Ash commented: We observe the onset of NADH oxidation at a bare pyrolytic graphite edge electrode at approximately +300 mV vs. the standard hydrogen electrode (SHE), *i.e. ca* +60 mV vs. Ag/AgCl. A range of other carbon electrodes have been demonstrated as NADH-based biosensors, and tend to report similar operating potentials as those stated for Prof. Abruña's graphene.¹

NADH oxidation to NAD⁺ can be summarised as involving the transfer of two e⁻s and one H⁺. The overall process has been proposed to follow an ECE mechanism involving at least one radical intermediate after removal of the first electron. In the case of biological NADH-dependent enzymes, side reactions involving any such radical intermediate or intermediates are avoided by rapid proton transfer and removal of the second electron. In the biological case, NADH oxidation is often coupled to highly regio- and stereospecific reduction of a substrate, with NADH effectively acting as a hydride donor that is required stoichiometrically during the course of the reaction and must be either regenerated or replenished.² Therefore, any process that leads to formation of biologically 'inactive' forms of NAD⁺ (for example due to dimerisation or reactions leading to immobilisation on an electrode) must be avoided in any practical NADH biosensor or NAD⁺ regeneration system. With this in mind it is surprising that you see no decrease in activity towards NADH oxidation with repeated potential cycles at your graphene electrode. It is possible that this is due to NADH depletion in the diffusion layer during your measurements limiting the overall extent of the reaction at the highest applied potentials, rather than the graphene electrodes being intrinsically inert to either reaction with or adsorption of NADH. The difference in activity between the mediator-coated and bare graphene surfaces could be that the adsorbed mediator limits the extent of NADH adsorption and reactive immobilisation, or that the semi-labile mediator (as indicated by the stated decrease in surface coverage with successive potential cycles) transfers two electrons simultaneously (or at least in quick succession) thus avoiding any long-lived radical intermediates.

You quote an overpotential of roughly 1 V for NADH oxidation – $E(\text{NAD}^+/\text{NADH})$ at pH 7 is approximately –400 mV vs. SHE at a 1000-fold excess of NADH, *i.e. ca* –640 mV vs. Ag/AgCl. These operating potentials are quite high in biological terms. Do you know what the activities of the bare or modified graphene electrodes are with respect to other possible substrates that might be present in a realistic biosensor? NADH-based biosensors sometimes incorporate a NAD⁺-dependent enzyme that produces NADH during substrate oxidation, and it is oxidation of this NADH by-product that is detected by the biosensor. It is therefore very important that the electrode oxidises exclusively NADH at the operating potential of the biosensor. An alternative strategy might be to use a more selective electrocatalyst that works at a lower overpotential to improve the selectivity of a NADH-based biosensor by minimising any unwanted oxidations.

- 1 See, for example: *Chapter 4: NAD(P)-Based Biosensors in Bioelectrochemistry: Fundamentals, Experimental Techniques and Applications*, ed. P. N. Bartlett, John Wiley and Sons, 2008.
- 2 Reeve *et al.*, *Chem. Commun.*, 2012, **48**(10), 1589–1591.

Elizabeth Howe asked: The graphene used in your study was used after being transferred with PMMA. How do you verify whether all the PMMA has been removed by your washing procedure?

Héctor Abruña responded: As mentioned in the manuscript, in order to remove the PMMA, the chips were immersed in anisole for two to three hours, in 50/50 (v/v) dichloromethane–acetone for at least 12 hours, and finally in isopropanol for at least two hours. The quality of the single layer graphene was verified by the 2D and G peaks at 2700 and 1600 cm^{-1} , respectively, in the Raman spectra. In some cases we also obtained AFM and TEM images that showed very low levels of PMMA residues.

Richard McCreery enquired: You described some Raman results in your manuscript. Does the Raman spectrum of your graphene change at all when you modify the graphene with polymers or other reagents?

Héctor Abruña responded: We see changes that we can relate to the nature of the films being deposited. In addition, the intensities of the bands associated with graphene itself (*i.e.* G and 2D) tend to be attenuated.

Richard McCreery continued: The ratio of the 2D peak intensity to that of the G peak in Raman is often used to indicate graphene structural changes. Does this ratio remain reasonably constant when your graphene is modified with polymers, etc?

Héctor Abruña answered: While we have not done an exhaustive/extensive study of this, we generally find that while there is some attenuation of intensities, the ratio remains relatively similar.

Katherine Holt said: A paper last year¹ reported that the hydrophobicity/hydrophilicity of graphene is strongly affected by the adsorption of atmospheric contaminants like hydrocarbons, which adsorb almost instantaneously on exposure to the laboratory atmosphere. Is this something that has been noted by those working on graphene electrochemistry and is it likely to effect the performance of graphene in electrochemical devices?

1 Z. Li *et al.*, Effect of airborne contaminants on the wettability of supported graphene and graphite, *Nature Materials*, 2013, **12**, 925.

Héctor Abruña replied: We have not carried out studies addressing this point. However, we would estimate that if there was strong adsorption of such molecules over a significant fraction of the surface, the redox response of redox molecules would be affected. We have not observed such behavior.¹

1 N. L. Ritzert, J. Rodríguez-López, C. Tan and H. D. Abruña, *Langmuir*, 2013, **29**, 1683–1694.

Julie Macpherson opened the discussion of the paper by Philippe Bergonzo: Fig. 7 of your paper shows that for your heavily doped diamond (1% doping) there is an oxygen wave present, occurring about -0.8 V in the planar surface CV. This is commonly associated with the presence of non diamond carbon (NDC) impurities. For your proposed applications does NDC presence matter?

Philippe Bergonzo replied: No, I don't think NDC carbon impurities matter.

Julie Macpherson continued: When you grow the heavily doped diamond films on structured surfaces, such as the CNT study described in the paper, does the NDC content increase compared to growth on planar surfaces – *i.e.* is it harder to control BDD quality on structured surfaces? It was difficult to tell from the data presented in Fig. 7 as the background currents had also increased due to an increased surface area of the electrode.

Philippe Bergonzo responded: According to the thickness of diamond deposited on the CNT structure, we observe a varying window, probably due to the (inverse) variation in NDC content.¹

¹ This is further discussed in C. Hébert *et al.*, *Carbon*, 2014, **71**, 27–33.

Richard McCreery remarked: Do electrochemical properties, particularly electron transfer kinetics, track the doping density? Has there been a systematic study of k^0 as a function of boron content in BDD?

Philippe Bergonzo replied: I don't know if such a systematic study has been conducted on all forms of diamond. For the material we grow, we indeed have published such a study.¹

¹ E. Vanhove *et al.*, High reactivity and stability of diamond electrodes: The influence of the B-doping concentration, *Physica Status Solidi a*, 2009, **206**, 2063–2069.

Richard McCreery said: You showed three reversible redox systems on boron doped diamond electrodes. One might expect a variable density of states in the BDD, possibly resulting in the dependence of kinetics on the standard potential of the redox system. Has anyone studied the variations in electrode kinetics with the E^0 of the redox system?

Philippe Bergonzo responded: Not to my knowledge. It could be interesting. But again one has to do it on well controlled materials before they can generalize this to 'diamond'.

Julie Macpherson commented: This is a comment related to the previous question Rick McCreery asked about the relationship between boron content, sp^2 content and the density of states in diamond. When boron doped diamond (BDD) is doped above the metallic threshold $>1 \times 10^{20}$ B atoms cm^{-3} , we find that the intrinsic rate constant of electron transfer, k^0 , for simple outer sphere redox species, scales with the boron content of the material, which in turn, we find, controls the local density of states (LDOS).¹ This was elucidated on polycrystalline BDD, which contains differently doped crystallographic grains and negligible sp^2 carbon. All grains were doped above the metallic threshold. A factor of ~ 4 difference in the measured k^0 values across the electrode surface was determined using scanning electrochemical microscopy (SECM), where more highly doped grains showed faster k^0 values (for example, for $FcTMA^+$ we measure k^0 values of 0.097 ± 0.04 $cm\ s^{-1}$ on high doped grains and 0.022 ± 0.008 $cm\ s^{-1}$ on lower doped ones); these k^0 values are high enough for CVs at macroelectrodes in stationary solution to look reversible, at typical scan rates *e.g.* 50 $mV\ s^{-1}$. Scanning electrochemical cell microscopy (SECCM) was used to measure the local capacitance in the differently doped regions,

from which a LDOS value could be extracted. The more highly doped regions showed a factor of ~ 4 increase in LDOS compared to the less doped areas; values of $1.7 \pm 0.7 \times 10^{20} \text{ cm}^{-3} \text{ eV}^{-1}$ and $6.3 \pm 2.0 \times 10^{20} \text{ cm}^{-3} \text{ eV}^{-1}$ were obtained for the polycrystalline BDD (at 0.0 V vs. Ag/AgCl) employed in the study. Notably, our 'metal-like' BDD had a LDOS 2–3 orders of magnitude smaller than for a metal, which is typically *ca.* $10^{23} \text{ cm}^{-3} \text{ eV}^{-1}$. Note, when sp^2 carbon is present, its affect on the LDOS, the electrical properties and the k^0 values should also be considered. The presence of sp^2 carbon should also be taken into consideration when interpreting the k^0 values for inner sphere couples, as sp^2 regions will have different surface chemistry with respect to the BDD surface.

1 This work was published in Patten *et al.*, *Angew Chemie Intl.*, 2012, 51, 7002–7006.

John Foord remarked: You showed this rather nice *in situ* pulsed electrochemical cleaning process for diamond electrodes. Have you tried it out on other materials where electrode fouling can be more problematic?

Philippe Bergonzo responded: We tried it on Pt, and for deposited thin Pt layers it seems to work, but the layers are progressively altered, something that is not observed on diamond. I suspect such a cleaning process may also work on bulk Pt materials or on other highly 'bulky' materials.

Keith Stevenson asked: You discussed the concept of self-cleaning of the electrodes and I was curious as to the potential dependence of self cleaning mechanism. Did you employ symmetrical pulses during the redox process?

George Zheng Chen also asked: When the electrode was cleaned by application of short cathodic and/or anodic pulses, were the pulses symmetrical? What potentials were selected, and would these potentials affect the cleaning result?

Philippe Bergonzo answered: Indeed they were symmetrical cathodic and anodic. We don't drive a voltage, but a current, typically from 10 to 100 mA maximum. More information is available in the following patent.¹

1 J. de Sanoit, R. Kiran and E. Scorsone, *US Pat.* 2013313120, 2013.

Keith Stevenson continued: As a follow up question, do you see a change in the behavior with a change in chloride ion concentration? Are bubbles observed due chlorine gas generation?

Philippe Bergonzo responded: We never saw any modification of the solution during the cleaning process. There are no bubbles. We believe this is the case because we go too fast; bubbles don't have time to form.

Patrick Unwin enquired: What does the voltammetry look like at the individual electrodes in these microelectrode arrays?

Philippe Bergonzo replied: We have published this.¹ We did a statistical study of how many of the microelectrodes were not responding as perfect diamond, and

indeed always found close to 15% faulty microelectrodes. We incriminated metallic residues from the final lithographic processing step, as this seemed to be improved after better cleaning. Thus it's not due to the diamond itself but to processing contamination. A minor issue as better cleaning can solve it.

1 R. Kiran *et al.*, *Sensors*, 2012, **12**, 7669–7681.

Julie Macpherson remarked: Given that sp^2 carbon is present in your highly doped boron doped diamond films, do you see any evidence for preferential etching of the sp^2 regions (*e.g.* at grain boundaries) when you apply extreme potentials?

Philippe Bergonzo answered: We don't use potentials but currents. And indeed their could be high surface potentials but I don't know how to evaluate them. What we have observed is that we have never seen any etching. Some colleagues (*e.g.* C.E. Nebel at IAF Freiburg) have observed diamond etching applying several 10s of V (tip -20 or -30 V). But we have never observed such an etching at the resulting 'potentials' of our activation treatments. For example, depositing IrO_x mediating spheres on the nanocrystalline diamond surface, we have observed closely how they behave, following their populations and sizes grain by grain, but we have never observed any etching on the diamond.

Katherine Holt said: It seems to be the case that sometimes the advantages of certain electrode materials (like diamond) need to be compromised to get the best out of them in a device. For example in the applications of diamond electrodes requiring interfacing with cells, is it not the case that the cells may not interact or grow well at the interface of the best quality planar diamond? So the diamond needs to be nanostructured, which introduces sp^2 content and defects, and therefore its electrochemical properties (wide potential window etc.) might be compromised?

Philippe Bergonzo responded: Nanocrystalline diamond is heavily structured at the nanoscale, and we have studied a lot of cell lines from neurons to organ cells or osteoblastes or cardiomyocytes, and they all grow extremely well on nanocrystalline diamond. In other words, I don't think we grow the "best quality planar diamond", and indeed we believe it works best without.

Katherine Holt continued: Does the surface termination of the diamond have any effect on the ability of cells to attach to it or the electrochemical response?

Philippe Bergonzo answered: Indeed. But not quite on the attachment. However, we have observed that O $-$ or H $-$ termination has a effect on the promotion or the reduction of glial to neuron cell growth. As such, we in fact deposited a patent application on this as it is a nice way to pattern the neuron growth on carbon materials with respect to localizing glial cells that are also required to build a neural network. This patent has been published recently,¹ where more details are provided.

1 S. Picaud *et al.*, *Eur. Pat.* 2747158 (A1), 2014.

Richard McCreery asked: Diamond electrodes are generally recognized as having longer lifetimes *in vivo*, probably due to lower adsorption of 'impurities', *etc.* But biological redox systems such as dopamine may also require chemical interactions with the electrode surface during electrocatalysis. Does the weak tendency to adsorb work against the need for catalysis? Does the stability of diamond become a disadvantage for the electrochemical detection of biological molecules?

Philippe Bergonzo responded: I don't think so, as electron transfer rates are still very high, reaching values close to 0.1 cm s^{-1} . However, this is only obtained on very highly doped diamond, typically in the 1% at. ratio range, thus it is likely that we are talking about diamond with a bit of sp^2 carbon.

Patrick Unwin commented: We have recently compared the voltammetric response for dopamine electro-oxidation on a variety of carbon electrodes.¹ There is a clear difference in the response between freshly cleaved HOPG (different grades) and boron doped diamond (BDD), with the process apparently much faster on HOPG and other graphite electrodes than on BDD. The comparison is interesting as HOPG and BDD have a broadly similar density of electronic states (DOS), yet the faster process on the graphite basal surface suggests stronger adsorption and electronic coupling. Such comparisons between graphite basal surfaces and diamond would be beneficial for other systems, given the similar DOS, but different structure and surface terminations (sp^2 vs. sp^3). It should be noted that although the electro-oxidation process is slower on BDD, the BDD surface is much less prone to blocking by dopamine oxidation products at relatively high dopamine concentration ($100 \mu\text{M}$), although at lower concentration ($1 \mu\text{M}$) HOPG shows minimal blocking and a fast response that is clearly seen due to the low background currents characteristic of pristine HOPG.

1 A. N. Patel, S. Tan, T. S. Miller, J. V. Macpherson and P. R. Unwin, *Anal. Chem.*, 2013, **85**(24), 11755–11764

Richard McCreery asked: In your experience with diamond electrodes for analysis *in vivo*, are biological redox systems detectable, given the often weak activity of diamond surfaces for the electrocatalysis involved in biological redox systems?

Julie Macpherson replied: Yes, biological redox systems are detectable, but given the electrocatalytic inert nature of the sp^3 diamond electrode surface, the detection potentials are typically pushed out compared to other carbon electrodes such as glassy carbon, carbon nanotubes, HOPG (EPPG, BPG) and classical metal electrodes such as Pt and Au, as we show, for example, in ref. 1 and 2. However, if this is not perceived to be a problem for the measurement of interest, diamond electrodes are well-suited towards *in vivo* measurements given their reduced fouling, bio-compatibility and mechanical robustness.

1 Patel *et al.*, *Anal. Chem.*, 2013, **85**, 11755–11764.

2 A. Güell *et al.*, *PCCP*, 2010, **12**, 10108–10114.

Matteo Duca said: Considering the interest in hybrid sp^2 – sp^3 carbon systems, do you believe that it is possible to control the BDD coating to prepare a partially-covered BDD–CNT composite using the approach described in your paper (Fig. 8)?

Philippe Bergonzo answered: Probably. This was partially discussed in ref. 1. However we always aimed at fully covering (no pinhole) diamond layers. When the layers are too thin, we see that the window decreases, and thus the advantages of the diamond layer for neurointerfacing as a passivating layer are lost. But for other applications it could be interesting indeed.

1 C. Hébert *et al.*, *Carbon*, 2014, **71**, 27–33.

Jet-Sing Lee asked: How much more effective are these boron doped diamonds compared to conventional materials for biosensing? I respect that conventional materials are more prone to fouling; are your materials using diamond stable enough to make them cost efficient for its enhanced properties?

Philippe Bergonzo responded: Boron doped diamond is costless. It is as cheap as any carbon based material to fabricate. We grow diamond from methane and water. The cost is definitely not an issue. Indeed we have shown that diamond is extremely stable *in vivo*, more than any other conventional carbon materials, and it is also much cheaper than any precious metal.

Richard McCreery opened the discussion of the paper by Dirk Guldi: One would expect a significant difference in the photoactivity of 'peapod' materials made with semiconducting compared to metallic nanotubes. How do you select the tubes used for making 'peapod' structures, *i.e.* what sort of pre-screening is carried out?

Dirk Guldi replied: There are indeed significant differences in terms of behaviour between semiconducting and metallic carbon nanotubes and between empty and filled carbon nanotubes. No attempts were, however, made to pre-screen the samples or to separate the different types of carbon nanotubes. Still, to guarantee unambiguous results we took great care to individualise carbon nanotubes and to ensure the stabilities of the corresponding dispersions. In addition, enriched semiconducting carbon nanotubes were used.

Yury Gogotsi asked: What is the percentage/degree of filling of the nanotubes by fullerenes and how does it affect the adsorption?

Dirk Guldi responded: The degree of filling is nearly 100%. In addition, filled carbon nanotubes, that is, peapods, require a slightly different work-up procedure than empty carbon nanotubes. The samples were quantitatively checked for filling *vs.* non filling by TEM.

Katherine Holt remarked: In your difference absorption spectra (for example Fig. 7 in the paper) you see very clear loss peaks (minima) corresponding to loss of the ground state absorptions. However only in a few cases do you see any corresponding gain peaks (maxima) in the difference spectra. Why is this?

Dirk Guldi answered: In general, the negative regions in the differential absorption spectra correlate to a loss of oscillator strength relative to the reference spectrum, that is, the ground state absorption. Overall, the spectra are correlated

with those obtained in spectroelectrochemical reduction and oxidation experiments performed with filled and non filled carbon nanotubes. One may expect, for example, upon one electron reduction, a bleaching of the ground state absorption accompanied by a transient absorption red-shifted relative to the latter. The fact that this is not always observable relates to the fact that we are dealing with a rather broad distribution of carbon nanotubes and, in turn, a superimposition of many spectra. As such, the bleaching seems in most cases more intense than the new absorption.

Katherine Holt asked: For the proposed applications for your functionalised carbon nanotubes would you require them to be dispersed in solution or deposited as a film?

Dirk Guldi responded: Ideal for an application would be carbon nanotubes deposited as films – 2D or 3D. One way to use the deposition would be to employ them in field effect transistors, in which the individual carbon nanotubes are electronically and spatially well isolated from each other.

Katherine Holt continued: To clarify the role of charge transfer in the dispersion of the functionalised nanotubes – if you tried to separate the nanotubes with a surfactant for which charge transfer does not take place between surfactant molecule and the nanotube, you will be unable to disperse the CNT?

Dirk Guldi replied: Just the bare separation is even possible for surfactants, which lack any redox features. However, the stability of the resulting dispersions is governed by increasing the activation barrier for the desorption of the surfactants from the surface of the carbon nanotubes. In other words, the redox features of the surfactants/charge transfer between the surfactants and the carbon nanotubes assist in raising the activation barrier.

François Béguin opened the discussion of the paper by Milo Shaffer: Your system is characterized by low power and energy density. The low energy density is related to the electrode construction with a very small amount of active material. Would it be possible to include activated carbon cloth in the structural composites? Activated carbon cloths from phenolic resin or viscose are currently available; these materials display a very high specific surface area and have good mechanical properties. Therefore, it is possible to believe that they could be substituted with a part of the structural fibers without greatly altering the mechanical properties. Could that be a future direction to overcome the low performance of your system? The low power density is essentially due to the low electrical conductivity of ionic liquid electrolyte. Which kind of alternative electrolyte could be considered to improve this parameter?

Milo Shaffer responded: Traditional activated carbon fibres are not suitable for structural supercapacitors as the intrinsic fibre properties are poor. They are not based on structural graphitised carbon, and have low strength, stiffness, as well as very poor electrical conductivity. They have a tendency to fragment during handling, increasing the risk of device shorts. Whilst they could be hybridised with conventional structural fibres, the reduction in structural fibre content

would damage mechanical performance. The aerogel approach is the most attractive, since it uses all of the matrix space to provide an electrochemically active electrode volume. The structural fibre content can therefore be fully maintained, whilst redeploying the remaining space that is always present, so a second function can be integrated. In due course, other fibres may become available that offer good strength and surface area simultaneously. Aligned carbon nanotube fibres are already promising in this context, and although not yet available in commercial quantity, efforts are underway to scale-up production.¹ The relatively low performance of the structural capacitors in the current study is due to the very poor conductivity of the basic cross-linked electrolyte used. Already we have produced a very significantly improved structural electrolyte based on a bicontinuous epoxy/ionic liquid system, as outlined in the paper, and in more detail elsewhere.² Once integrated with the aerogel/carbon fibre electrodes, significant performance improvements can be anticipated.

1 N. Behabtu *et al.*, *Science*, 2013, **339**(6116), 182–186.

2 N. Shirshova *et al.*, *J. Mat. Chem. A*, 2013, **48**, 15300.

Yury Gogotsi said: It is important to think of applications when deciding whether to use a multifunctional energy-storing composite or a supercapacitor film on a strong structural composite plate.

Milo Shaffer replied: As indicated in the paper, the combination of two fully independent mechanical and electrical systems by simple adhesion, to form a multifunctional structure, may offer some advantages particularly in usable system volume. However, the scope for weight saving is very limited, on the whole relating only to packaging. Conceptually, it is also a much less interesting or revealing challenge. Multifunctional materials, in which at least some of the atoms provide at least two traditionally orthogonal functions, are more demanding but potentially much more rewarding. The motivation to balance conflicting materials requirements may inspire new understanding and new materials solutions of broad relevance.

Yury Gogotsi continued: What is better in terms of strength and capacitance per weight – using your multifunctional composite or a strong structural composite combined with a thin film supercapacitor?

Milo Shaffer responded: As noted in the previous answer, the multifunctional materials approach has much greater potential to save both weight and volume. The materials challenges are, of course, also greater.

Richard McCreery asked: You see a very large increase in capacitance (approx. 100×) when the carbon aerogel is added to other carbon materials. Is this all due to the increase in surface area, or is there also a component due to formation of oxygen functional groups? The latter are important to activated carbon, for example.

Milo Shaffer answered: In the examples within the paper, the capacitance of the electrodes closely follows the BET surface area measurements; the as grown carbon nanotubes form part of this series, even though their synthesis does not

involve intentional oxygen-containing precursors. Some oxygen groups will be present but only at a low level. It is also worth noting that these examples have mostly mesoporosity accessible to the electrolyte. The consistency between surface area and these various electrodes suggests that surface groups contribute little, in this case. Deliberate attempts to increase the effective capacitance by introducing pseudocapacitive species is of interest, as illustrated in the paper by the polymerisation of PEDOT onto the aerogel framework, although there are a wide range of possibilities.

George Zheng Chen commented: Regarding the influence of the oxygen content in nano-structured carbon materials on their specific capacitance, I shall mention that when presented in an appropriate concentration in surface functional groups, oxygen does contribute to the so-called pseudo-capacitance that involves electron transfer reactions. The cyclic voltammograms of such materials do not necessarily present current peaks because these oxygen containing surface groups are connected by highly conjugated graphene layers, which means that the electrons transferred to and from these groups are also likely partly or fully delocalised and hence do not give rise to current peaks. Our experience is that the oxygen content should not be higher than 10% to benefit the specific capacitance. Too high an oxygen content may disrupt the surface conjugation of the carbon material. You may find more explanations on the relation between pseudo-capacitance and delocalised valance electrons in my recent paper (which is free to access on the journal's webpage).¹

1 G. Z. Chen, Understanding supercapacitors based on nano-hybrid materials with interfacial conjugation, *Prog. Nat. Sci.-Mater. Int.*, 2013, **23**(3), 245–255.

Richard McCreery enquired: Oxygen functional groups are often responsible for the degradation of high surface area carbon materials. Since the aerogels are made from oxygen-rich precursors, are they prone to corrosion more than other carbon materials?

Milo Shaffer answered: It is true that RF aerogels are formed by the condensation of oxygen-containing precursors; however, many other activated carbons are similarly derived from carbohydrate sources, usually with other heteroatoms present. The oxygen content can be adjusted by the carbonisation temperature, which for our aerogels is around 800 °C. Other cross-linking chemistries could be introduced to the same concept/architecture if necessary, though the RF system is considered favourable for supercapacitor electrodes.¹

1 A.G. Pandolfo and A.F. Hollenkamp, *J. Power Sour.*, 2006, **157**, 11–27.

Richard McCreery continued: Has anyone tried H-termination to stabilize carbon aerogels? It is not a simple process, but should remove oxygen surface groups and make stable C–H bonds.

Milo Shaffer replied: Controlled termination of the surface chemistry is an important process for controlling electrolyte interactions, and the associated effects on electrolyte wetting and stability.¹ In the context of aerogels, thermal activation or passivation with oxides have been considered.² Introducing

hydrogen terminations would be worth exploring for these monolithic aerogels. We have explored high temperature activated radical reactions³ that might be applied, without the complications of plasma systems.

1 A. G. Pandolfo and A. F. Hollenkamp, *J. Power Sour.*, 2006, **157**, 11–27.

2 J. Biener *et al.*, *Energy Environ. Sci.*, 2011, **4**, 656–667.

3 R. Menzel *et al.*, *Chem. Sci.*, 2010, **1**, 603–608.

Thomas Rabbow asked: Did you try to combine first a chemical activation to achieve a high active surface area and second the aerogel coating to achieve a high total surface area? Thereby I would not expect a higher BET surface area, but if the chemical activation would persist for a long enough duration, I could imagine that the capacitor performance increases (faster charging and discharging, lower internal resistance).

Milo Shaffer answered: We did try various combinations of the strategies for improving the structural electrode surface area/capacitance, which are outlined in the paper. As yet, we have found no significant advantage. Introducing the RF after an initial fibre activation is likely to block the surface pores created in the first step. Activation after the aerogel formation ought to offer some advantage, but relatively less than might be anticipated from the increased BET surface area, due to the formation of inaccessible pores.¹

1 A. G. Pandolfo and A. F. Hollenkamp, *J. Power Sour.*, 2006, **157**, 11–27.

Fulian Qiu commented: When a double layer supercapacitor is fully charged, its open circuit potential will gradually decay due to self discharge processes. These processes could include a localized short circuit through the separator membrane, thermo-desorption, back reaction at the substrates, reactions involving impurities on carbon particles and surface morphology changes due to strain–stress under the charged conditions.

Fulian Qiu asked: What are the advantages of composite supercapacitors over the activated carbon coated metal mesh supercapacitor?

Milo Shaffer replied: The point is that the primary carbon fibres are truly structural, of the type used for a wide range of structures under high mechanical load – for example, airframes, chassis, sports equipment, *etc.* Thus, in principle, the composite can be truly multifunctional – both a mechanical structural component and an electrical energy storage device.

George Zheng Chen remarked: Would you please comment on the differences between your laminated structural supercapacitors and the cable-type structural supercapacitors that can be woven or embedded into different structures?

Milo Shaffer responded: There are a wide range of possible architectures for structural supercapacitors, with the cells implemented at the single fibre, tow, or laminate level. In general, using a finer scale design offers improved rate performance at the expense of complicating the connections and increasing the risk of shorts. Twisted yarn supercapacitors and batteries have been produced

previously, with a focus on flexible systems, for example, to integrate in textiles. In our structural supercapacitors, the structural requirements are more demanding; good stiffness, in particular, demands good fibre alignment, and is less consistent with existing cable approaches. On the other hand, structural supercapacitor architectures designed around predominantly tensile loading of unidirectional fibres would be attractive. The preparation of multifunctional matrices which balance both ionic conductivity and stiffness are challenging, implying that matrix dominated failures are likely to be limiting. Purely tensile loading in the fibre direction should offer a simpler route to good multifunctional performance.

Héctor Abruña said: The processes that you described are likely to significantly add to the cost of the materials and thus the devices. In the long run, will the improved performance compensate the additional costs?

Milo Shaffer replied: Of course, the balance of cost to benefit will ultimately determine the take up in commercial application. However, it is an early stage to predict either the costs of any finalised system or the extent of the benefits available. Neither the materials systems nor processes are settled. We are at the stage of a proof of concept, exploring the fundamental constraints and compromises involved. Different applications will have very different requirements. The value of weight saving will vary dramatically between, for example, the aerospace and automotive sectors. However, as an indication, even in the cost-sensitive mass produced car market, there is intense interest in purely structural carbon fibre composites to replace metals, particularly for hybrid and electric vehicles. Once carbon fibres are implemented, the marginal cost of structural supercapacitors may be modest. There may even be savings, since a secondary electrical energy system will be redundant. The work has also generated a new type of hierarchical composite which has wider potential significance than structural capacitors alone, since there are indications of improved absolute mechanical performance and a range of other opportunities.

Elzbieta Frackowiak asked: In your experiments concerning multifunctional composites for a rigid energy storage device, you were using a very narrow range of potential, ± 0.2 V, that is quite far from the voltage applied for a capacitor operating in practice. What is the maximum voltage which you reached in the two solutions used, *i.e.*, KCL and ILs?

Additionally, could you please answer what is the benefit of using a step voltage (0.1 V) method for capacitance estimation? Are you sure that capacitance, ESR and BET surface area values (Table 1 and 2) are measured with so high accuracy?

Milo Shaffer responded: Broadly, our test regime was driven by the challenges of making a true structural supercapacitor. Many of the samples, indeed all the structural ones, in the paper still have a high ESR and low power density, for reasons explained in the text. In order to test all samples under consistent conditions in a reasonable timeframe, we adapted our approach, as you mention. Once we have plausibly practical values, which we are approaching with our next generation architecture, then of course it will be important to test under a range of more realistic operating conditions. For non-aqueous electrolytes, where we hope to broaden the voltage window significantly, a new challenge will arise: how

to process and seal the composite such that it remains sufficiently dry. However, we have plans and commercial collaborators in place to address this requirement.

Boris Dyatkin remarked: Your publication describes two approaches of grafting nanotubes onto carbon fibers. The first involves catalytic growth, with the carbon nanotubes stemming directly from the carbon structure. The second involves the decomposition and carbonization of a solvent containing suspended nanotubes. Does the latter approach increase charge transfer resistance in the system, given that the electrical contact between the two carbon structures is not an ideal, defect-free graphitic interconnect? What is the structure of the carbonized residue that forms the interconnect?

Milo Shaffer replied: We have not explored this question in detail, due to the modest capacitances identified, but found no evidence of such a problem. The high temperature of the carbonisation, the very small distances involved, and the network forming ability of the nanotubes are likely to limit the influence of the amorphous carbon resistance. Note that in the sizing (second) approach, the nanotubes lie predominantly parallel to the carbon fibre surface, due to capillarity effects, implying a large number of contact points per nanotube.

Jet-Sing Lee asked: I understand that for your material its structure is important for applications where weight is important. You show a capacitance of $\sim 15 \text{ F g}^{-1}$ for your CAG material. Is this not still lower than reported activated carbons that can reach up to $\sim 200 \text{ F g}^{-1}$, therefore less can be used? You say an aim of your material is to address mechanical issues but from my knowledge, carbon and other materials that operate by electric-double layer capacitance are already robust with stabilities of up to 500,000 cycles.

You also mentioned that you are now modifying your CAG with conducting polymers to enhance its performance. How stable are these new composites, as polymeric supercapacitors that work *via* pseudocapacitance are known to have low stabilities?

Milo Shaffer answered: The capacitance mentioned only appears lower because it is normalised to the combined mass of the carbon fibre and CAG (hybrid) material. The capacitance of the CAG normalised to its own mass is within the usual range for these materials. The carbon fibres contribute to current collection, but little to the capacitance. The concept uses the necessary matrix volume within the composite to accommodate CAG and hence electrical energy storage, at the same time as maintaining the structural performance of the primary fibres.

The pseudocapacitive components have the same advantages and limitations as in conventional supercapacitors, but can contribute to the overall performance of this novel multifunctional system. The motivation, here, is not at all the simple mechanical robustness (stability) of a pure electrochemical electrode. The point is that these multifunctional composite devices may potentially have sufficient mechanical performance to be truly structural, in the sense of carrying the load needed, for example, to act as a wing spar, car door, or aerofoil.

Robert Dryfe opened the discussion of the paper by George Chen: Please could you clarify the motivation behind the study. If the objective is electrochemical

energy storage, the utility of the approach seems limited because the efficiency (based on the anodic/cathodic charge ratios given in Table 2) seems low. Repeated cycling would yield <1% charge recovery within 10 cycles.

George Zheng Chen responded: The data in Table 2 are derived from cyclic voltammograms recorded at different potential scan rates, and show clearly that the Q^+/Q^- ratio increases with decreasing scan rate. This suggests some kinetic difficulties. In the text following Table 2, we describe that if the potential is held at the cathodic end for a short while (*e.g.* 10 s), the Q^+/Q^- ratio can increase to close to 1. In practice, the charging and discharging are more likely carried out under a steady state, and we can expect improvement in charge efficiency upon optimisation of process variables. In fact, Table 3 contains a Q^+/Q^- ratio of 0.83 when the deposition and dissolution were carried out under a constant current of 150 mA. In addition to fundamental understanding, there are however engineering issues in terms of what type of electrode/cell designs will be optimal to avoid the deposited carbon falling off the electrode when the deposit becomes thicker. Other factors, such as formation of CO gas, may also reduce the charge efficiency. I shall point out that for seasonal (and long term) energy storage, carbon presents many merits over other options, particularly its high energy density and convenience for transport and storage. Cost is another important factor affecting commercial development which favours the molten salt approach, not to mention that there is the tax benefit for capture and utilisation of CO₂.

Robert Dryfe continued: Following on from the previous question, do you have any idea of where the ‘missing’ charge is going, *i.e.* what are the side reactions that lead to the loss of charge recovery?

George Zheng Chen answered: The side reactions include the formation of CO gas, and possibly alkali metals, both of which may leave the cathode and are not recoverable. The deposited carbon, if sufficiently thick, may also fall off the cathode, accounting partly for the missing charge.

Matteo Duca said: Since you observe that “carbon deposition occurs on the deposited carbon more efficiently than on a freshly prepared metal electrode” (line 35, page 11 of your paper), have you tried using graphite as a cathode instead of metals?

George Zheng Chen replied: We have tested using graphite as the cathode material, but it suffered from lithium intercalation which agrees with what was reported previously.^{1,2}

1. G. Z. Chen *et al.*, *High Temp. Mater. Processes*, 1998, 2, 459–469.

2. G. Z. Chen *et al.*, *J. Electroanal. Chem.*, 1998, 446, 1–6.

Yury Gogotsi asked: Is there any fundamental reason why you can't use a lower temperature for electrochemical CO₂ conversion?

George Zheng Chen responded: There are indeed fundamental reasons. From thermodynamic analysis, electrochemical reduction of CO₂ should proceed at room temperature at a fairly low voltage ($\Delta G = 394.363 \text{ kJ mol}^{-1} \Rightarrow \Delta E = 1.022 \text{ V}$ at 273 K).

However, the mechanistic and kinetic factors play more important roles here. It is difficult to find a molecular or ionic acceptor that can work in low temperature electrolytes to capture and stabilise the O^{2-} ion that is discharged from direct reduction of CO_2 , and then transport the O^{2-} ion for discharging at the anode.

Nevertheless, I shall mention that working at elevated temperatures for electrolysis is not a big technical or energy concern because the industry is familiar with using molten salts at 1173–1273 K for the extraction of aluminium which has always been a profitable business.

Martin Jonsson-Niedziolka enquired: When calculating efficiency, what part of the energy consumption goes into maintaining the high temperature needed for the process?

George Zheng Chen explained: Energy consumption for supplying the electrolysis current is far greater than that for heating the molten salt to maintain the working temperature. Heat loss is obviously inevitable, but can be minimised *via* the use of appropriate thermal insulation materials and also reactor design. Particularly, the Joule heat resulting from the electrolysis current passing through the molten salt should be considered in balancing the heat loss. In our laboratory, the thermal energy loss from our furnace, which is far from optimal for heat insulation, has been estimated to be about 5% or less of the total energy consumption. On the other hand, I shall mention that an important prospect of our work is to utilise electricity and heat that can be generated from sunlight or other renewable sources, or that from the excess generation of existing power plants.

Manuel Alvarez-Guerra commented: I think these types of approaches for CO_2 valorisation can be particularly interesting as a future option for using excess energy from intermittent renewable sources (*e.g.* solar or wind energy), allowing the balancing of fluctuations in production and demand cycles, which is usually considered to be essential for intermittent renewable energies to be used on a large scale.

Manuel Alvarez-Guerra asked: In your paper you have some results of the electrolysis under different CO_2 partial pressures using mild steel as cathode. Did you find important changes in the performance and deposition rates depending on the nature of the working electrode used?

George Zheng Chen answered: In the early stage of carbon deposition, it depends strongly on the nature of the working electrode, except for graphite which suffers from lithium intercalation. However, after the electrode surface is fully covered by carbon, the deposition become very much independent of the substrate materials. Also, we have not found notable differences between samples obtained from using different cathode materials, but we have not analysed the carbon layer immediately next to the surface of the cathode.

Manuel Alvarez-Guerra continued: You carried out the experiments in a two electrode undivided cell. Did you try using a divided cell with a membrane?

George Zheng Chen replied: Thank you for your question. The use of a membrane or separator in our cell has been considered, but this option increases the cell resistance and hence energy consumption. It also remains challenging in terms of membrane materials. Ideally, the membrane should be made from an ionic conductor that is stable in molten salts at elevated temperatures. However, all known high temperature ion conducting materials (*e.g.* zirconium stabilised yttrium oxide or sodium beta alumina) are ceramic types and vulnerable to thermal shock and molten salt attack. Porous ceramic materials are an alternative, but these would be less effective in terms of separation of the anolyte and anodic products from the catholyte and cathodic product. However, there are still engineering innovations that can help improve the process efficiency and product quality, and these are under investigation in our laboratory and possibly others.

Héctor Abruña remarked: Data comes from the automotive industry like Ford and GM.

George Zheng Chen answered: Thank you. I have checked a few recent review articles¹ on PEMFCs, and it seems that in some laboratories, the power density has reached beyond 1 W per sq. cm, which implies a current density of over 2 A per sq. cm at a cell voltage of 0.5 V.

1 *e.g.* H. Zhang and P.K. Shen, *Chem. Rev.*, 2012, **112**, 2780–2832.

Domnik Bayer said: Your paper presenting a direct carbon energy storage is appealing. This implies the reversible reduction of CO₂ and re-oxidation of deposited carbon at a reasonable efficiency. My question relates to the system cost in terms of power specific (€ per kW installed) and energy specific costs (€ per kWh stored). These numbers would make comparison with existing storage solutions easier (*e.g.* battery technologies). Carbon production from CO₂ might also be interesting, however, if there are no additional special features of the electrochemically generated carbon, it will pricewise compete against standard carbon black which is available at a sales price of ~1000 US\$ per metric ton.

George Zheng Chen responded: We have not yet constructed a full device (*e.g.* direct carbon fuel cell) to test and evaluate the energy efficiency of the carbon in great confidence. For discussion, we have estimated the cost of our electrolytic carbon to be \$3–5 per kg-C, which includes the electricity, heat and post-electrolysis processing costs based on our laboratory findings, but not yet the benefit from the carbon tax. Also, for energy storage, the cost for post-electrolysis processing may not be needed. With further research and technical improvement, the process efficiency can increase further, and the cost may be driven below \$3 per kg-C. The energy content in pure carbon is about 9.2 kWh per kg-C (enthalpy change of carbon combustion at 600 °C), whilst the direct carbon fuel cell can achieve over 70% in energy efficiency, leading to a delivery of 6.4 kWh per kg-C. Thus, \$0.5 per kWh in energy stored is not impossible. For material uses, there are many options such as electrode materials for supercapacitors and batteries. Specifically, the molten salt approach reduces the carbon footprint, whilst all existing processes for production of various carbon materials strain the environment with some levels of carbon emission.

François Béguin stated: Simply based on the energy consumption of electrolysis, you estimate the cost of your carbon to be less than \$5 per kg, and you compare it to the current market price of activated carbon for supercapacitors (\$20 to 40 per kg). The relatively high cost of supercapacitor carbons is due to the post-treatments required to reach a high purity. Therefore, the estimated cost of your material should include all the washing treatments which are necessary to reach the same purity as for the supercapacitor carbons.

George Zheng Chen responded: I fully appreciate your comment. Our estimation has been based on the preparation procedures and materials/chemicals used in our laboratory, including washing with acid, which is of course subject to correction when the process is scaled up for commercial purposes. Our intention to report our preliminary estimate is to provide a starting point for more comprehensive cost analyses so that potential stakeholders, researchers and business managers alike, can have a scale for comparison.

Chi-Chang Hu said: Is it possible to control the microstructures (*e.g.* porosity, crystallinity, doping, *etc.*) of carbon by means of electrochemical deposition?

George Zheng Chen replied: Your question is also mine. According to our results obtained so far, there is some degree of controllability *via* manipulating the temperature and voltage (or the cathode potential) as we have described in ref. 1. I hope our ongoing research can provide more accurate answers, at least in terms of the controlled formation of a particular structure.

1 H. V. Ijije, C. Sun and G. Z. Chen, *Carbon*, 2014, 73, 163–174.

Laurence Hardwick asked: Have you investigated what type or form of carbon you produce electrochemically from CO₂ – say for example with Raman?

George Zheng Chen answered: We have carried out some basic spectroscopic and microscopic characterisations of the electro-deposited carbon, and reported the findings in ref. 11 of our paper. However, we have not yet analysed our samples by Raman, although you may find ref. 3 and 10 of the paper to be relevant.

Milo Shaffer remarked: What is the structure of the carbon deposit? Does its integrity, structure, or electrical conductivity/contact resistance limit the thickness of coating that might be deposited effectively? From the images that you showed, it appears that the deposit is porous, containing either voids due to gas evolution, or trapped electrolyte. If graphitic carbons form, is (undesirable) intercalation favoured?

George Zheng Chen responded: The deposited carbon layer is fairly porous and formed by packed nano-particulates of carbon in various morphologies. The mechanism for these nano-particulates to be fairly stably attached to the electrode surface is not clear, and deserves further investigation. My initial explanation is that it is related to the insolubility of carbon in the molten salt. Thus, once formed, the carbon nano-particulates would precipitate to the nearest solid

Discussions

surface such as that of the cathode. It is not to say that carbon deposition would not produce a crystalline structure *via* the nucleation and growth mechanism, but the fast deposition rate, by product (CO gas and alkali metal) formation and gas bubbling on the nearby anode, could all disturb the growth of a crystalline deposit. The contact resistance between these nano-particulates, if any, is inevitable, but is also very well compensated by the increase of the overall surface area when more carbon is deposited, and hence does not affect the growth of the deposit. As for intercalation, it will surely happen to graphitic carbon, as we know it is possible to produce carbon nanotubes by intercalation of lithium or sodium into graphite in molten salts.¹ Although intercalation may not occur to amorphous carbon or the carbon nano-particulates, alkali carbides may still form, although this was found to be insignificant.

¹ *cf.* (a) Chen *et al.*, *High Temp. Mater. Processes*, 1998, 2, 459–469; (b) Chen *et al.*, *J. Electroanal. Chem.*, 1998, 446, 1–6.