CHAPTER 4

Microfluidic Reactors for Nanomaterial Synthesis

S. Krishnadasan, A. Yashina, A.J. deMello and **J.C. deMello***

Contents

Abstract

The difficulty of preparing nanomaterials in a controlled, reproducible manner is a key obstacle to the proper exploitation of many nanoscale phenomena. An automated chemical reactor capable of producing (on demand and at the point of need) high-quality nanomaterials, with optimized physicochemical properties, would find numerous applications in nanoscale science and technology, especially in the areas of photonics, optoelectronics, bio-analysis, and targeted drug-delivery. In addition such a device would find immediate and important applications in toxicology, where it is essential to characterize the physiological effects of nanoparticles not only in terms of chemical composition but also in terms of size, shape, and surface functionalization. In this chapter, we describe recent advances in the development of microfluidic reactors for

Department of Chemistry, Imperial College London, South Kensington, London SW7 2AY, UK.

* Corresponding author:
E-mail address: j.demello@imperial.ac.uk

Advances in Chemical Engineering, Volume 38
ISSN: 0065-2377, DOI 10.1016/S0065-2377(10)38004-5

controlled nanoparticle synthesis and, more specifically, work in our group aimed at developing just such an automated reactor.

1. INTRODUCTION

The discovery in recent years of novel properties, processes, and phenomena at the nanoscale has created revolutionary opportunities for the creation of novel materials and devices with superior chemical, physical, and/or biological characteristics (Ozin and Arsenault, 2005). Nanocrystalline materials are of particular interest in this regard owing to their tunable physico-chemical properties and their potential use as functional elements for biological sensing, optoelectronics, fiber-optic communications, and lasers (Alivisatos, 1996; Euliss et al., 2006; Green, 2004; Han et al., 2007; Kawazoe et al., 2006; Klostranec and Chan, 2006; Matsui, 2005; Medina et al., 2007; Rhyner et al., 2006). The characteristics of nanocrystals are strongly influenced by their physical dimensions, and there is consequently considerable interest in processing routes that yield nanoparticles of well-defined size and shape (Donega et al., 2003).

There are two main routes to nanoparticle formation: (1) "top-down" and (2) "bottom-up" approaches. In top-down routes, nanometer-sized structures are engineered from bulk materials using a combination of lithography, micromachining methods, and etching (Xia et al., 2007). Significantly, the creation of sub-100-nm structures requires lithographic techniques beyond the optical domain, such as electron beam and X-ray lithography. Such approaches are technically challenging and although reproducible do not readily lend themselves to large-scale production. A bottom-up approach on the other hand involves the chemical growth of particles on an atom-by-atom or molecule-by-molecule basis until the desired particle size and shape are achieved (Malik et al., 2005). This growth process occurs spontaneously in super-saturated solutions and has been successfully used to create high-quality spherical, cubic, tubular, and tetrahedral crystallites in kilogram quantities and above (Masala and Seshadri, 2004; Milliron et al., 2004; Park et al., 2004). The bottom-up approach, which may be carried out on the laboratory bench using standard techniques in synthetic chemistry, has attracted considerable interest owing to its versatility and ease of use, and it is by far the most practical and prevalent means of producing large quantities of nanomaterials. In recent years, a variety of sophisticated chemical strategies have been reported for producing near defect-free nanoparticles of consistent size, shape, and chemical composition. Their implementation, however, remains a complex and difficult undertaking that requires a combination of skill (and often luck), intuition, and extensive experimentation to obtain well-defined nanoparticles with tightly specified properties (Donega et al.,

2003; Dushkin et al., 2000; Rao et al., 2006; Yordanov et al., 2006). Indeed the formidable difficulty of preparing high-quality nanoparticles in a controlled and reproducible manner is recognized to be the foremost obstacle to the full exploitation of many nanoscale phenomena. In this chapter, we consider the utility of microfluidic systems in providing a controlled environment in which to synthesize nanomaterials and also the feasibility of automating the synthesis procedure, with a view to creating "blackbox" chemical reactors that in response to appropriate instructions—and without any human intervention—can produce high-quality nanomaterials with optimized physicochemical properties. The provision of such automated tools would significantly advance the field of nanoscience, allowing high-quality nanomaterials to be produced on-demand and at the point of need for a host of applications in nanotechnology.

2. MICROFLUIDIC ROUTES TO THE SYNTHESIS OF NANOPARTICLES

The approach described later on in this chapter builds upon a report in 2002, in which we proposed microfluidic reactors as favourable systems for nanoparticle synthesis, and showed that nanocrystalline cadmium sulfide prepared in such reactors exhibited improved monodispersity compared with particles prepared in conventional bulk-scale vessels (Edel et al., 2002).

In simple terms nanoparticles are formed via a two-stage process in which an initial nucleation stage (in which seed particles spontaneously precipitate from solution) is followed by a more gradual growth phase in which diffusion of solutes from the solution to the seed surface proceeds until the final particle size is attained. This classical model was first proposed by LaMer and Dinegar to explain the mechanism of formation of sulfur colloids (La Mer and Denegar, 1950) and can be represented by a simple cartoon of the kind shown in Figure 1 which shows the variation in solute concentration as a function of time. The solute is formed by a chemical reaction (e.g., hydrolysis of a metal alkoxide, hydration of metal ions, and decomposition of organic compounds). As the reaction proceeds, the solute concentration will at some point exceed a "supersaturation" concentration, eventually reaching a "critical" concentration at which point nucleation occurs (ideally) in a short burst. This nucleation process—and the ensuing growth of these nuclei—lowers the solute concentration to a value below the critical nucleation concentration, thereby preventing further nucleation, but still allowing particle growth. The formed particles then grow at a rate that merely consumes all further solutes that are generated by the chemical reaction. The process of particle growth will lower the overall free energy of the system so, in the absence of any other competing processes, growth will

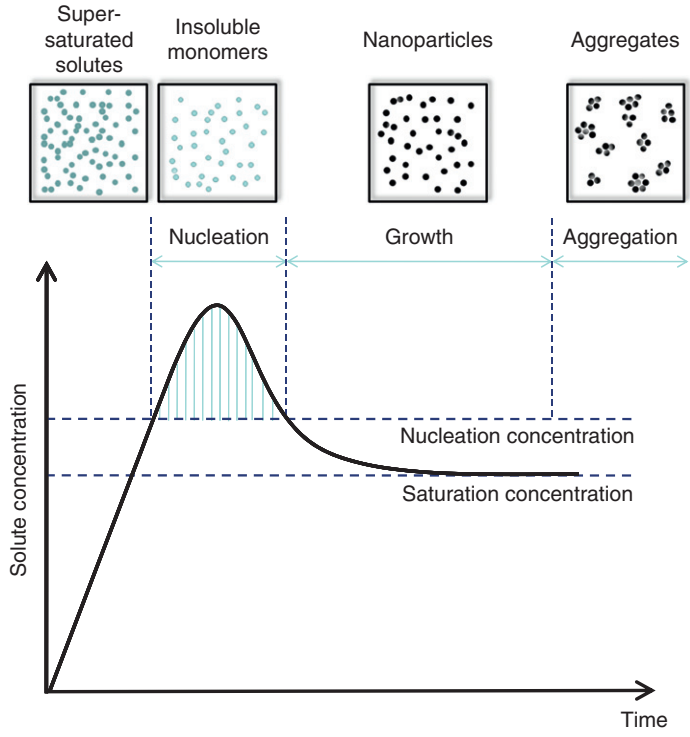

Figure 1 Cartoon illustration of nucleation and growth during the preparation of monodisperse nanoparticles.

continue until all of the solute has been consumed. Additionally, aggregation of individual particles also lowers the free energy of the system so the particles will tend to coalesce over time and precipitate out of solution. Unfortunately, in most cases, nucleation and growth occur concurrently, and thus the final particle population will therefore exhibit a broad (and undesirable) size distribution. Accordingly, to obtain monodisperse particles, it is necessary to set up conditions in which all nucleation takes place over a short period of time with additional material being supplied so gradually that it find its way to the nuclei without the solute concentration reaching a level at which further nucleation can take place. In practical terms it is crucial that all nuclei should form and grow in an identical environment with state functions (such as temperature, pressure, and concentration) assuming constant values throughout the reaction volume. In conventional syntheses within bulk reactors (where turbulent mixing is used to ensure rapid combination of reagents) significant variations in physical conditions across the reaction chamber are typical, thus generating wide particle-size distributions. In the current context, microfluidic systems—which allow for rapid and controlled thermal and mass

transfer—are an ideal format for nanoparticle production. They are able to control the temperature or temperature gradient along a flow profile and can rapidly heat or cool the reagent mixture. They can efficiently mix reagents on sub-millisecond time scales and allow reagents to be added in a flexible and controlled manner.

In subsequent studies by ourselves and others, a variety of microfluidic architectures have been used to prepare metal and compound semiconductors nanoparticles, including CdS, CdSe, TiO_2, Ag, Au, and Co (Boleininger et al., 2006; Chan et al., 2003, 2005; Cottam et al., 2007; DeMello and DeMello, 2004; Khan et al., 2004; Kohler et al., 2005; Krishnadasan et al., 2004; Millman et al., 2005; Shalom et al., 2007; Shestopalov et al., 2004; Song et al., 2006; Takagi et al., 2004; Wagner and Kohler, 2005; Wang et al., 2004, 2005; Xue et al., 2004; Yen et al., 2003, 2005). In simple terms, microfluidic devices manipulate and process sub-microliter volumes of liquid in enclosed channels that are typically no more than a few hundred microns in diameter (see Figure 2). Importantly for chemical synthesis, they have a number of advantageous features, including precise control over reaction conditions, rapid heating, cooling and mixing of fluid streams, and the ability to combine multiple chemical processes into a single, integrated device (deMello, 2006). Over the years, numerous functional elements have been successfully integrated into microfluidic devices (including filters, heaters, mixers, valves, actuators, electrophoretic separators, distillers, and various kinds of transducer), and it is now possible to perform virtually every conceivable chemical process using appropriate combinations of these elements. This has given rise to the concept of a lab-on-a-chip, a self-contained integrated device that is capable of performing all steps in a complete chemical synthesis or analysis without any additional equipment or instrumentation.

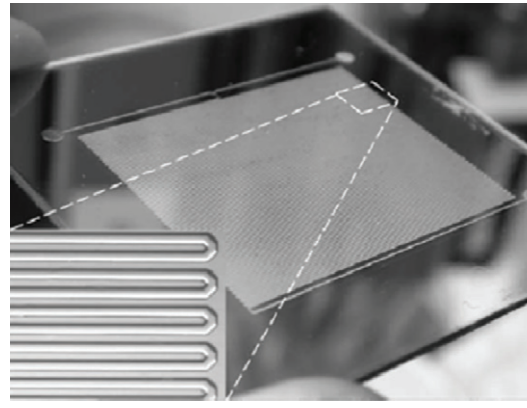

Figure 2 A simple microfluidic device fabricated in glass using wet lithography. The typical channel dimensions in a microfluidic device are 5–100 μm.

In the early days, microfluidic devices were used mainly for chemical analysis due to their intrinsic ability to handle small sample volumes and the significant performance enhancements that typically arise when analytical techniques are transferred to the microscale (Manz et al., 1992; Vilkner et al., 2004). However, in recent years, there has been increasing interest in using them for chemical synthesis where the precise control they provide over reaction conditions offers potential improvements in product yield, purity, and quality (deMello, 2006). Microfluidic devices can be fabricated using a variety of substrate materials, including glass, silicon, and polymers such as poly(methyl-methacrylate) and poly-(dimethylsiloxane) (Dumais et al., 2006; Iliescu, 2006). Of these materials, glass and quartz are generally preferred for nanoparticle synthesis due to their transparency in the visible region of the spectrum and tolerance of elevated reaction temperatures. Glass devices are normally prepared using wet chemical etching (see reference Iliescu, 2006), but other techniques such as direct laser-etching procedures have also been used successfully. Devices of considerable sophistication have been reported in the literature, as shown, for example, in Figure 3, but the benefits of miniaturization can also be exploited in surprisingly simple chip architectures as will be seen below.

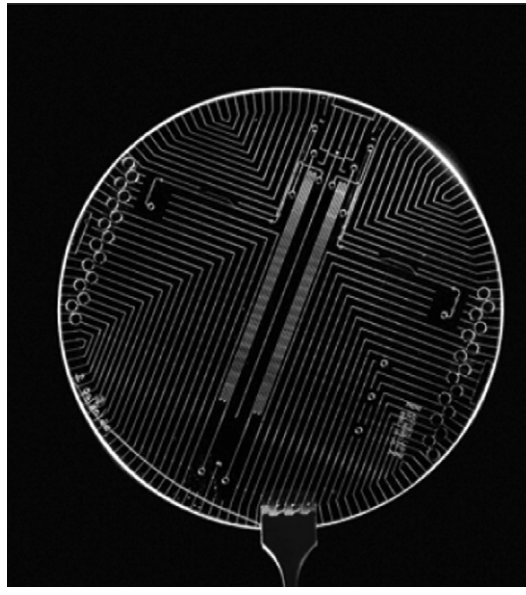

Figure 3 A sophisticated microfluidic device used for DNA processing. Numerous functional elements have been adapted for use with microfluidics, and it is now possible to perform virtually any combination of chemical processes in a microfluidic device.

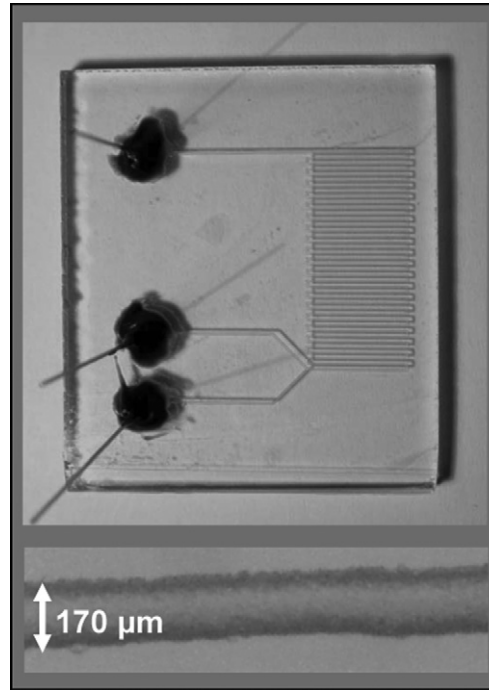

Figure 4 A simple y-shaped microfluidic device with two inlet channels and a single outlet channel. The device can be used to mix and react two reagents A and B. By varying the relative volumetric rates (F_A and F_B) at which the reagents are injected, it is possible to vary the composition of the mixture in the outlet channel. By varying the total flow rate ($F = F_A + F_B$), it is possible to control the time the reagents spend in the reaction channel.

Microfluidic reactors comprise one or more inlet ports into which reagents are typically injected using precision syringe pumps and capillary connectors, a network of channels in which various chemical processes are carried out, and one or more outlet ports where products and waste materials are extracted. Figure 4 for instance shows a simple y-shaped chemical reactor with two inlet channels and a single outlet channel that can be used to mix and react two reagents A and B. The composition of the resultant mixture and the residence time inside the microfluidic chip can be controlled by varying the volumetric flow rates of the two reagent solutions, F_A and F_B. Hence, if the volume of the outlet channel is V and the molar concentrations of the two reagent solutions are [A] and [B], then the residence time (τ) in the reaction outlet channel is given by

$$\tau = \frac{V}{F_A + F_B} \tag{1}$$

and the molar ratio (R) of A to B in the final mixture is equal to.

$$R = \frac{F_A}{F_B} \times \frac{[A]}{[B]} \tag{2}$$

In a typical scenario, the volume of the outlet channel might be 10 µl and the individual injection rates of the two reagent solutions might each be of order 100 µl min^{-1}. Using modern syringe pumps, it is straightforward to control each of the injection rates to within 0.1 µl min^{-1}, and so in these circumstances τ and R can each be controlled to within 0.1%, which far exceeds the precision achievable in conventional macroscale reactors. In the case of thermally initiated reactions, the y-shaped reactor can be placed on a hot-plate with high spatial uniformity. In this way, the entire reaction volume can be held at a uniform temperature that can be controlled to within a fraction of a degree, which similarly exceeds the typical temperature uniformity achievable in a conventional macroscale reactor.

An additional advantage of using microfluidic devices, which we do not have the space to discuss in detail here, is the absence of turbulence (Koo and Kleinstreuer, 2003). In the context of nanoparticle synthesis, turbulence gives rise to unpredictable variations in physical conditions inside the reactor that can influence the nature of the chemical product and in particular affect the size, shape, and chemical composition. In microfluidic devices, turbulence is suppressed (due to the dominance of viscous over inertial forces) and fluid streams mix by diffusion only. This leads to a more reproducible reaction environment that may in principle allow for improved size and shape control.

These improvements in control do not come without cost since the transferral of chemical reactions from the macroscale to the microscale frequently requires compromises to be made in terms of reagent selection and reaction conditions. The solvents used for microfluidic syntheses, for instance, should ideally be low-viscosity liquids at room temperature (mp << 20 °C) as this conveniently avoids the need to heat the syringe pumps or to thermally lag the capillaries that connect the syringe pumps to the inlets of the microfluidic device. And, in the case of standard lyothermal syntheses, the solvents should also have sufficiently high boiling points to permit thermal initiation of the chemical reactions (bp >> 250 °C). The former requirement precludes the use of many of the standard solvents used in conventional macroscale synthesis routes (such as long-chain alkylamines) as these tend to be waxy materials with high melting points. Unfortunately, solvents with low melting points frequently have relatively low boiling points and so tend to be unsuitable for many lyothermal syntheses. The choice of solvent system for microfluidic-based syntheses is therefore somewhat restricted compared with macroscale syntheses. In addition, it is generally not possible to use the

extremely high reaction temperatures favored for many macroscale lyothermal syntheses ($T > 300\,°C$) since above about 280 °C glass chips have a tendency to fracture and the epoxy-based glues used to attach the capillaries to the microfluidic chips tend to decompose. The reactions must therefore be carried out at lower temperatures, which risks increasing the number of defects in the resultant nanocrystals. These issues, although mundane in nature, mean it is rarely possible to implement existing synthesis routes directly on a microfluidic chip without making at least some adaptations to the solvent, reagent mixtures, and/or reaction temperature.

In this chapter, we will focus on the automated preparation of CdSe quantum dots (which have been extensively studied and characterized in the scientific literature), but the principles and techniques we describe are general ones that may be applied to a variety of nanomaterials. Indeed, the following paragraphs outline some selected studies which have reported the use of microfluidic reactors for the synthesis of compound semiconductor nanomaterials.

Nakamura and coworkers have reported the synthesis of CdSe nanoparticles in simple glass capillaries (Nakamura et al., 2002). The authors utilized the ability to control both temperature and the mixing environment to rapidly synthesize CdSe particles from pre-mixed trioctylphosphine: selenium and cadmium acetate. Using such an approach, they were able to continuously produce CdSe nanoparticles at temperatures between 230 and 300 °C in 7–150 s. The size and quality of the obtained nanoparticles were assessed using absorbance spectroscopy and residence time distributions were reduced by introducing 500 nl nitrogen bubbles into the flow at defined intervals. Subsequently, the same group advanced the complexity of the synthesis process nanoparticles by generating CdSe/ZnS core/shell nanoparticles on the microscale (Wang et al., 2004). Specifically, preprepared reagents (CdSe, ZnS) were allowed to react within a chip-based microfluidic reactor, and two separate oil baths were used to maintain defined temperatures for the nucleation and growth of CdSe and ZnS, respectively (Figure 5a). CdSe nanoparticles with different ZnS shell thicknesses could be generated by simply varying the flow rate and capillary length. Luminescence spectra of uncoated CdSe and ZnS-coated particles are shown in Figure 5B, demonstrating the utility of the process. More recently the authors have refined their approach to demonstrate the synthesis of ZnS/CdSe/ZnS quantum-dot quantum-well (QDQW) structures (Uehara et al., 2009). In essence QDQWs are analogous to planar quantum-well devices where the narrow band gap inner shell acts as a quantum well embedded in a wide-gap material. By altering the size of the core and the thickness of both inner and outer shells, the electronic properties of the QDQW may be controlled. In many respects continuous flow microfluidic systems are ideal tools for synthesizing QDQWs since both

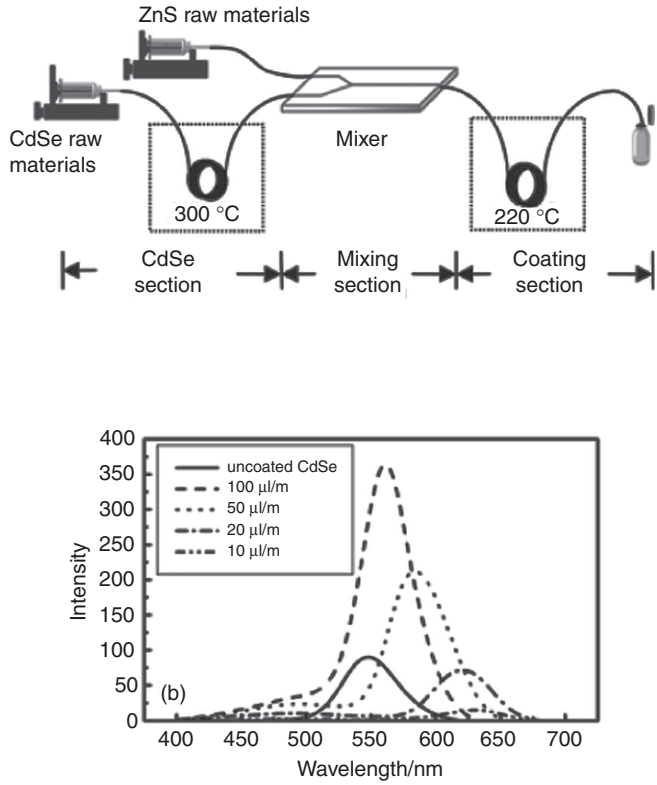

Figure 5 A schematic illustration of a multistep continuous-flow synthesis system for making ZnS-coated CdSe composite particles. Luminescence spectra of uncoated CdSe from CdSe portion and coated particles from the outlet at different flow rates. Labels denote the flow volume of the two syringes for the CdSe and ZnS raw feedstock. Uncoated CdSe nanoparticles were obtained using a flow rate of 100 μl min^{-1} images reproduced, with permission, from Wang et al., 2004).

the heating time and temperature can be controlled precisely. To demonstrate efficacy, the authors mixed a ZnS colloid solution with a core CdSe population in a continuous flow at 240 °C. The formed ZnS/CdSe (core/shell) solution was subsequently mixed with ZnS at 150 °C to produce the ZnS/CdSe/ZnS (core/shell/shell) multilayer composite nanoparticles. Such a high temperature synthesis is favorable since it leads to a high-quality crystalline coating, but requires careful control to avoid Ostwald ripening of cores that would lead to particle polydispersity. The authors importantly showed that the resulting ZnS/CdSe/ZnS QDQWs emit blue fluorescence with a fluorescence quantum efficiency as high as 50%. Moreover, photoluminescence could be tuned with ease by varying the volumetric flow rate during the CdSe deposition process.

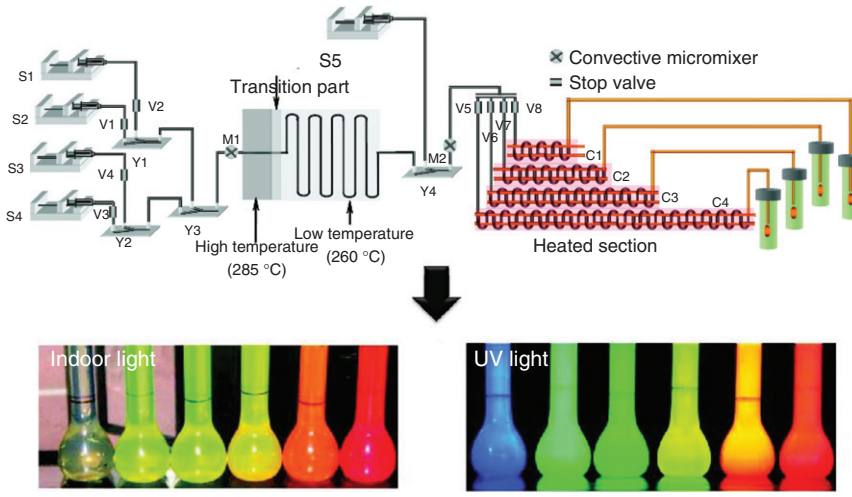

Figure 6 Schematic of an integrated microreactor for the continuous synthesis of CdSe/ZnS and CdS/ZnS nanoparticles (S1-syringe pump with Se precursor, S2-syringe pump with S precursor, S3-syringe pump with Cd-OA-OLA, S4-syringe pump with Cd-OA-OLA-TOPO, Y−Y conjunction, M-micromixer, V-stop valve, C-channel;images reproduced, with permission, from Yang et al., 2009).

More recently, Yang et al. have extended this idea and demonstrated the ability to perform multistep reactions in a sequential manner by producing ZnS-capped CdSe and CdS nanocrystals in a microfluidic device (Yang et al., 2009). This system involved distinct temperature zones for nucleation and growth of the core and a serpentine reaction channel for coating the ZnS shell (Figure 6). Such a system provides for both efficient mixing and uniform residence times for the precursors at high volumetric flow rates. The authors demonstrated the continuous synthesis of strongly emitting CdS/ZnS and CdSe/ZnS nanocrystals through real-time adjustment of precursor and capping inputs. Importantly, luminescence quantum efficiencies ranged from 41% for blue emitting particles and 78% for yellow particles.

A number of elegant studies over the past few years have also addressed the need to minimize particle size distributions through the use of segmented flow microfluidic systems. Such an approach removes the possibility of particle deposition on channels and eliminates the problems of residence time distributions that occur in single phase systems (where drag at the channel walls sets up a velocity distribution inside the channel). For example, Shestopalov et al. reported the two-step synthesis

of colloidal CdS and CdS/CdSe core–shell nanoparticles in a droplet-based microreactor (Shestopalov et al., 2004). In addition, Chan et al. have described the use of microfluidic droplet reactors for the high-temperature synthesis of CdSe nanoparticles (Chan et al., 2005). In this study, Cd/Se precursor solutions were made to form stable nanoliter-sized droplets in a perfluorinated polyether continuous phase. The encapsulated reagents were reacted when heated to 290 °C to yield nanoparticles 3.4 nm in diameter. In a similar study, Jensen et al. used gas–liquid (rather than liquid–liquid) segmented-flow reactors incorporating distinct temperature zones for the synthesis of high-quality CdSe quantum dots (Figure 7) (Yen et al., 2005). More recently, Hung et al. synthesized CdS nanoparticles in droplets by the passive fusion of the droplets containing different reagents ($Cd(NO_3)_2$ and NaS) (Hung et al., 2006). Controlled passive fusion was achieved using a dilating channel geometry and controlled liquid-phase flow. Particles obtained using this approach were generally smaller than those obtained by bulk and ranged between 8.2 and 4.2 nm in diameter. Significantly, all of the above studies leveraged efficient mixing and reduced residence-time distributions to engender improvements in both yield and size distribution. As an aside, it is noted that higher order nanostructures (inaccessible via conventional routes) can be synthesized using static microdroplet reactors. For example, Millman et al. reported the synthesis of anisotropic particles using nanoliter-sized droplets which are made to float on the surface of a perfluorinated oil (Millman et al., 2005). Since such floating droplets can be controlled by electrical fields, droplets containing suspensions of polymers and nanoparticles can be persuaded to form complex particle structures. Indeed, "striped" multilayer particles were generated from ternary mixtures of gold, fluorescent latex, and silica particles; and core–shell particles could be synthesized by the encapsulation of droplets of aqueous suspensions inside polymer shells.

Interestingly, the majority of compound semiconductor research to date has focused on II–VI systems such as CdSe due to the relative simplicity of the synthetic routes. III–V materials, whilst potentially less toxic, have presented a stiffer synthetic challenge. In large part, the strong covalency of the constituent atoms makes it hard to devise labile precursors, which leads to synthesized particles with permanent crystal defects. Very recently, Nightingale and deMello reported the first synthesis of III–V (InP) nanoparticles using a microfluidic reactor (Nightingale and de Mello, 2009). Using a simple microfluidic mixer (maintained at a temperature in excess of 200 °C) and product detection using a 355 nm diode-pumped Nd:YAG laser and a fiber-optic-coupled charge-coupled device (CCD) spectrometer, the authors verified that microfluidic reactors can be readily applied to III–V materials, yielding particles that are of comparable quality to those obtained using bulk methods.

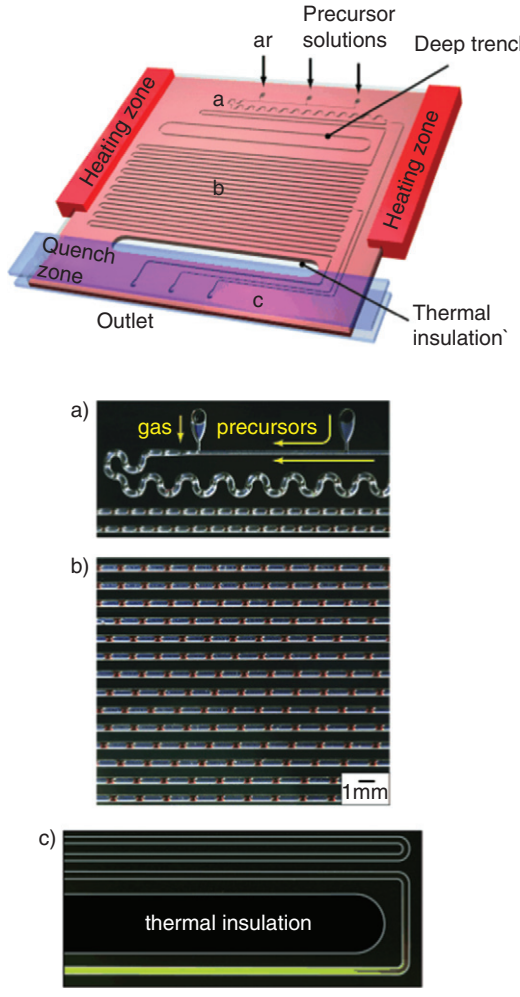

Figure 7 A microfluidic reactor for CdSe nanoparticle production. (A), The reactor provides for rapid precursor mixing (sector A), particle growth (sector B), and reaction quenching (sector C). The reactor accommodates a 1-m-long reaction channel and two side channels for collecting reaction aliquots. A halo etch region allows localization of temperature zones for reaction (>260 °C) and quenching (<70 °C). Precursor solutions are delivered into the heated section separately, and an argon gas stream generates a segmented gas–liquid flow. Recirculation within the liquid slugs rapidly mixes reagents and initiates the reaction. The reaction is stopped when the fluids enter the cooled outlet region. Photographs of heated inlets (B) and main channel section (C). Red segments show the reaction solution; dark segments define Ar gas; $T = 260$ °C; gas flow rate $= 60\ \mu l\ min^{-1}$; liquid flow rate $= 30\ \mu l\ min^{-1}$; images reproduced, with permission, from Yen et al. (2005).

3. THE AUTOMATED PRODUCTION OF CdSe NANOPARTICLES

The synthetic route we use in this work is a simple adaptation of a method by Peng et al., in which CdO and elemental Se are reacted together at high temperature in the presence of oleic acid to form CdSe nanocrystals (Peng and Peng, 2001). In brief, a precursor Se solution is prepared by combining 30 mg of Se with 10 ml of 1-octadecene and 0.4 ml of trioctylphosphine and warming over a hot plate. A Cd precursor solution is prepared by combining 13 mg of CdO and 0.6 ml of oleic acid in 10 ml of 1-octadecene and heating at 180 °C until clear. The reaction is performed in a glass y-shaped microfluidic chip with channels of width 330 μm and depth 160 μm. The reaction channel is 40 cm long and arranged in a serpentine architecture for compactness. The chip is placed on a stabilized hot plate with high spatial uniformity. The reaction can be performed at temperatures in the range 160–255 °C. Two syringe pumps are used to inject the precursor solutions into the inlet channels at rates up to 40 μl min^{-1}. The solutions mix rapidly at the point of confluence, and nucleation and growth of the CdSe nanoparticles occurs along the reaction channel. The emergent particles can then be monitored at ambient temperature at an observation zone downstream of the reaction zone. Here, the particles are excited using a 355-nm solid-state laser excitation source and emission is detected using a fiber-optic CCD spectrometer. A fraction of the incident laser light is redirected to a Si photodiode using a beam-splitter, allowing the emission spectra to be corrected for variations in the laser intensity. The set up is shown in Figure 8.

Figure 9a shows a TEM image of a typical CdSe nanoparticle prepared using the above synthesis. The emission spectrum of the nanoparticles obtained at an illustrative reaction temperature of 220 °C and equal

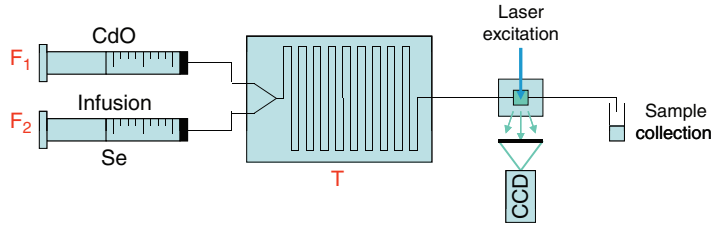

Figure 8 A schematic of the reactor used to synthesize the nanoparticles described in this chapter. Cd and Se precursor solutions are stored in two separate syringes and injected at flow rates $\mathbf{F_1}$ and $\mathbf{F_2}$ into the two inlets of a y-shaped microfluidic device. The microfluidic device rests on a hot plate of variable temperature ***T***. The reagent streams meet at the point of confluence and nucleation, and growth of the particles occurs as they pass along the outlet channel. The emission spectra of the particles so produced are monitored prior to collection at a detection-zone downstream of the chip using a 355-nm Nd : YAG laser as an excitation source and a fiber-optic-coupled CCD spectrometer.

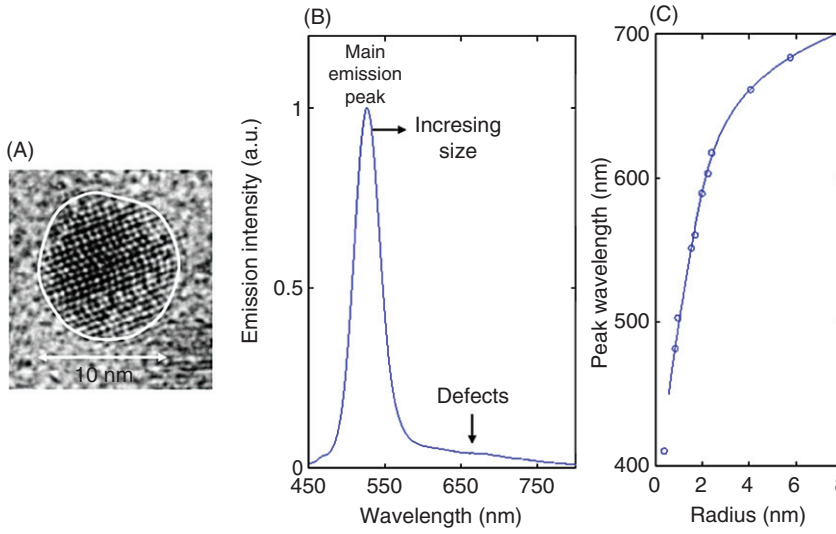

Figure 9 (A) Typical TEM image of a CdSe nanoparticle synthesized by the direct reaction of Se and CdO (see main text). (B) Typical emission spectrum of a CdSe nanoparticle. The spectrum exhibits two main features: (1) a sharp Gaussian-shaped peak due to band-edge emission and (2) a broad feature at lower energies due to defect emission. (C) The size-dependence of the peak wavelength for CdSe nanoparticles, determined using data provided in Murray et al. (1993).

CdO and Se flow rates of $9\,\mu l\,min^{-1}$ is shown in Figure 9b. The spectrum comprises a strong band-edge emission peak at 525 nm with full-width half-maximum 35 nm, and a broad weaker peak at 660 nm due to emission from crystal defects at the surface of the nanoparticles. Murray et al. have previously reported emission data for a series of size monodisperse nanoparticles, and it can be seen from their data (reproduced in Figure 9c) that the wavelength of the band-edge emission peak shifts progressively to longer wavelengths with increasing particle size, consistent with reduced quantum confinement effects.

The power of the microfluidic approach is apparent in Figures 10 and 11 where we consider the effects of varying T and τ. Increasing the temperature and extending the reaction time have similar effects on the band-edge emission, leading in both cases to a red-shift and enhancement in the intensity. This is consistent with the formation at higher temperatures and longer reaction times of larger particles, in which excitons are less tightly constrained and so less susceptible to trapping at surface defects (Pradhan et al., 2003). In both cases, a remarkable degree of control is achieved over the nanoparticle properties—with, for example, the peak wavelength being tunable to within a fraction of a nanometer—indicating the power of the microfluidic approach.

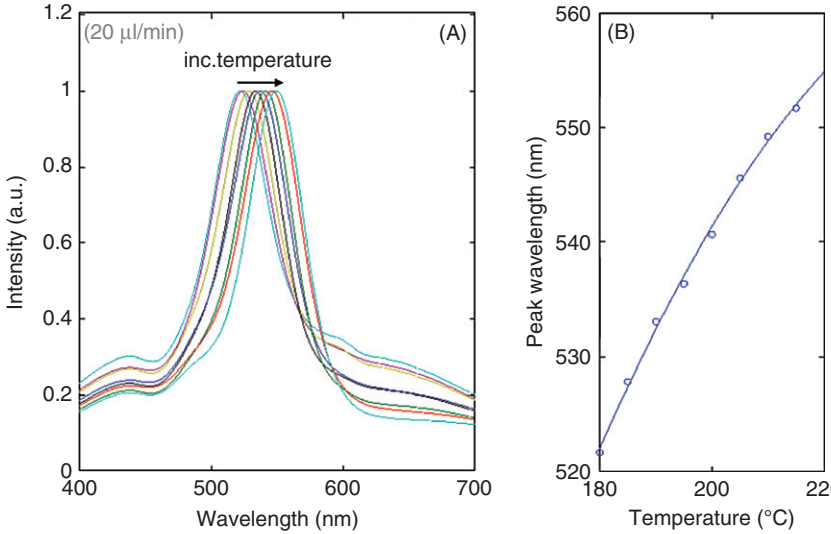

Figure 10 (A) Temperature dependence of the emission spectra of CdSe nanoparticles prepared at a fixed flow rate of 20 µl min^{-1} using the microfluidic system shown in Figure 8; (B) Temperature dependence of the peak wavelength, determined from the data in (A).

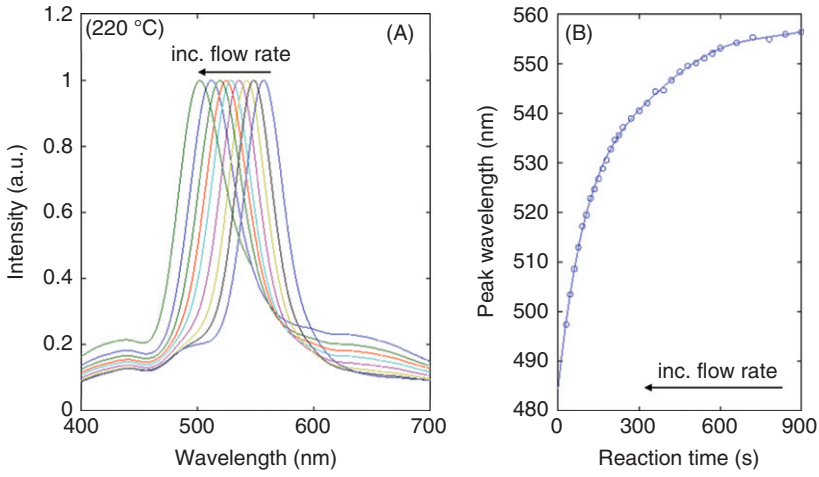

Figure 11 (A) Total flow-rate dependence of the emission spectra of CdSe nanoparticles prepared at a fixed temperature of 220 °C using the microfluidic system shown in Figure 8; (B) Total flow-rate dependence of the peak wavelength, determined from the data in (A).

4. THE AUTOMATED PRODUCTION OF NANOPARTICLES

As noted in the introduction, a major aim of the current research is the development of "black-box" automated reactors that can produce particles with desired physicochemical properties on demand and *without* any user intervention. In operation, an ideal reactor would behave in the manner of Figure 12. The user would first specify the required particle properties. The reactor would then evaluate multiple reaction conditions until it eventually identified an appropriate set of reaction conditions that yield particles with the specified properties, and it would then continue to produce particles with exactly these properties until instructed to stop. There are three essential parts to any automated system—(1) physical machinery to perform the process at hand, (2) online detectors for monitoring the output of the process, and (3) decision-making software that repeatedly updates the process parameters until a product with the desired properties is obtained. The effectiveness of the automation procedure is critically dependent on the performance of these three subsystems, each of which must satisfy a number of key criteria: the machinery should provide precise reproducible control of the physical process and should carry out the individual process steps as rapidly as possible to enable fast screening; the online detectors should provide real-time low-noise information about the end product; and the decision-making software should search for the optimal conditions in a way that is both parsimonious in terms of experimental measurements (in order to ensure a fast time-to-solution) and tolerant of noise in the experimental system.

The microfluidic system described above has two key features that make it especially amenable to creating such automated reactors: firstly, the reaction conditions can be precisely and rapidly manipulated (which in turn means that it is possible to finely tune the physical properties of the reaction product), and secondly, the inline spectrometer provides immediate real-time information about the product. The only remaining element

Figure 12 Schematic illustrating the desired behavior of an automated chemical reactor. The user enters the desired particle properties, the "black-box reactor then evaluates multiple reaction conditions until it identifies an appropriate set that yield particles with the desired properties; the reactor then continues to produce particles with these properties until instructed to stop.

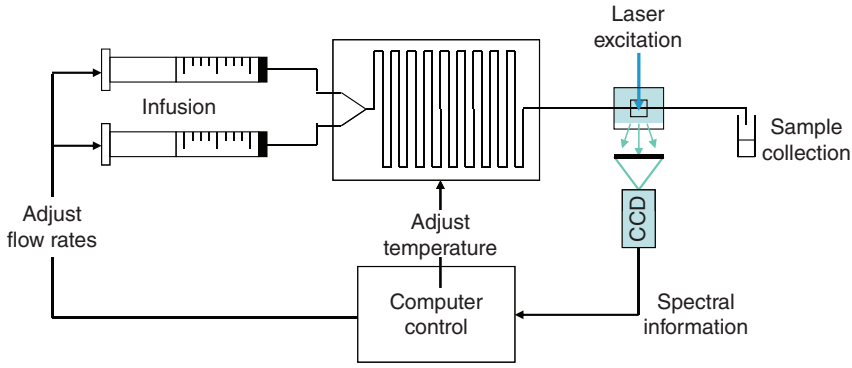

Figure 13 Schematic of an automated system for producing nanoparticles with desired properties. The set up is an adaptation of the system shown in Figure 8. The emission spectra of the emergent nanoparticles recorded by the CCD are passed to an intelligent control algorithm that repeatedly updates the reaction temperature and the injection rates of the two reagents until particles with the desired properties are obtained.

needed to create a fully automated reactor is the addition of a suitable control algorithm that is able to repeatedly update the reaction conditions until particles with the desired properties are obtained (see Figure 13).

The intention of the remaining part of this chapter is to explain in simple terms how the process of nanoparticle synthesis can be automated, and we will not dwell excessively on mathematical details. Instead, to illustrate our general approach, we will start by considering a very simple example in which we will design an automated system capable of producing particles that emit at a specified wavelength. We will then show how this approach can be generalized to enable the automated production of particles whose emission characteristics are specified in more complex ways. The key to automating the process of nanoparticle synthesis is to use a so-called *utility function* that reduces all of the known information about the particles to a single figure of merit that characterizes the particle quality. The utility function is usually defined in such a way that the figure of merit decreases steadily (from a large positive number) to zero as the measured properties get progressively closer to the desired properties. For example, let us suppose that one wishes to produce particles that emit at a certain target wavelength λ_t and that for the current reaction conditions the measured wavelength is λ_c. In this case, a sensible utility function $u(\lambda_c)$ would be

$$u(\lambda_c) = (\lambda_c - \lambda_t)^2 \tag{3}$$

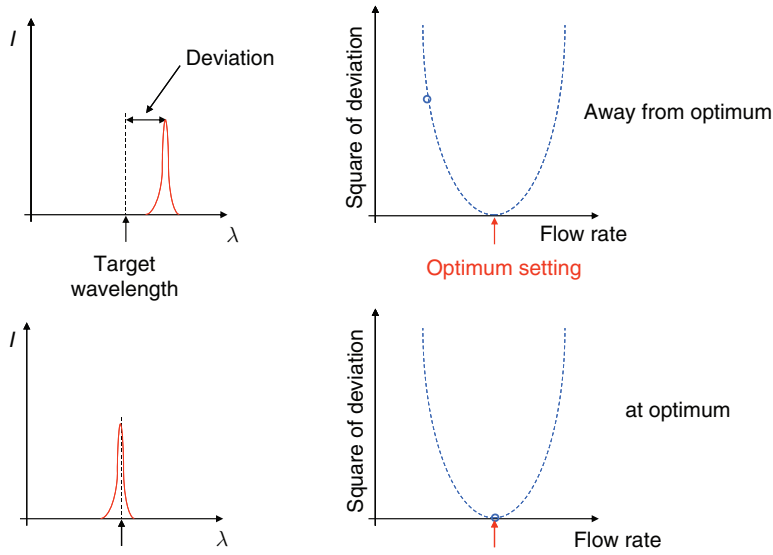

Figure 14 Schematic illustrating the use of a "utility function" to optimize the emission wavelength by varying the total flow rate of the reagents. The utility function generates a figure of merit that equals the square of the deviation of the current wavelength from the target wavelength. The square of the deviation is large and positive when the current wavelength is far from the target and decreases to zero as the target wavelength is approached. The optimization routine adjusts the flow rate in an effort to minimize the figure of merit and, in so doing, indirectly finds the reaction conditions that yield particles with the desired emission properties. The same approach can be readily extended to multiple reaction variables and the simultaneous optimization of multiple attributes (e.g., peak wavelength, peak intensity, and line-width).

which has a value of zero when the current wavelength equals the target wavelength and becomes progressively larger in size as the current wavelength deviates (in either direction) from the target.[1] The minimization routine tries to minimize the figure of merit and, in so doing, indirectly finds the reaction conditions that yield particles that emit with the target wavelength as shown in Figure 14. The precise manner in which the routine goes about finding the minimum is clearly crucial to the success of the optimization, and a number of issues need to be taken into account in choosing or designing a suitable routine:

1. The detailed mechanisms of nanoparticle formation (nucleation, growth, aggregation, and ripening) are understood only in qualitative detail so there are no reliable process models available

[1] Note, an alternative figure of merit would be the absolute deviation from the target wavelength. This choice of utility function, however, would exhibit an abrupt change of slope at the optimum which is liable to cause numerical difficulties for the control algorithm. The slope of the proposed parabolic utility function varies smoothly about the optimum and hence avoids these difficulties.

to guide the automation. This differs from the usual situation in chemical engineering where mathematical models of the chemical plant are normally available to supplement measured data. In the case of nanoparticle synthesis, the only information available to the algorithm is the measured data. The utility function in effect behaves like a black box that converts the process conditions to a final figure of merit in a manner that is entirely unknown to the minimization routine.

2. The same nominal reaction conditions may on different occasions give rise to slightly different products due to mechanical or chemical imperfections in the reactor. The utility function is therefore said to be noisy. The vast majority of minimization routines are designed for handling noise-free mathematical functions and only a very small subset of algorithms can cope with appreciable noise.
3. Owing to the slow nature of nanoparticle formation, there is typically a delay of several minutes between setting new reaction conditions and the figure of merit reaching a new stable value. The utility function is said to be expensive to evaluate and, from a practical perspective, this means only a small number of reaction conditions can be tested during the search for the optimum process conditions. This restricts the choice of algorithm severely since conventional minimization routines often require thousands of measurements to find the optimum conditions. If each measurement takes on average 5 min, one thousand measurements would correspond to a full week of searching, which is clearly excessive for many applications. Ideally, the control algorithm should be able to find the optimum reaction conditions in one hundred measurements or less.
4. The reaction conditions are constrained. In other words, there is usually a strict upper and lower limit for each reaction parameter. In the case of the synthesis described above, for example, the lower temperature is set by the need to provide sufficient thermal energy to initiate the reaction and the upper temperature by the need to remain below the decomposition temperature of the glue (see Section 2). The lower and upper limits on the total flow rate meanwhile are determined, respectively, by the maximum length of time one is prepared to allow for a single reaction and the minimum reaction time needed to produce crystals of nanometer dimensions. In this work, we select minimum and maximum total flow rates of 2 and 40 $\mu l\,min^{-1}$ which, for the typical chip volumes we use (~16.6 μl), correspond to average residence times of about 500 and 25 s, respectively.
5. The majority of minimization routines are designed for unconstrained optimization, in which the control algorithm is free to select any parameters it wishes. Only a minority can handle constrained optimization.

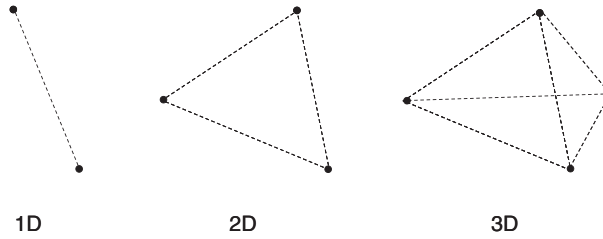

Figure 15 Schematic illustrating the concept of a simplex. A simplex is a geometric shape formed from $N+1$ vertices in an N-dimensional space. Hence, for one, two- and three-dimensional spaces, the simplex points comprise the vertices of a line, triangle, and tetrahedron, respectively. The simplices can move through their respective spaces by undergoing repeated reflections and/or changes of shape (see main text).

There are very few minimization routines that satisfy all of these conditions (i.e. that can perform constrained optimisation of expensive noisy black-box functions), and it is only in the past few years that effective algorithms have started to emerge. One promising technique, developed originally for noise-free optimization, is the simplex method (Kolda et al., 2003). A simplex is a geometric shape formed from $N+1$ distinct points (vertices) in an N-dimensional space. This is shown in Figure 15a–c for one-, two-, and three-dimensional spaces, where it can be seen that the simplex points comprise the vertices of a line, triangle, and tetrahedron, respectively. In the most basic form of simplex optimization, the value of the function is evaluated at each of the vertices and the simplex is repeatedly reflected away from the worst vertex. This is illustrated in Figure 16 for the illustrative two-dimensional function $f(x,y) = x^2 - 4x + y_2 - y - xy$. With each successive reflection, the simplex moves progressively closer to the minimum at (3,2). The simplex ceases to make progress at the 28th iteration (T28). The reason is straightforward. Vertex A has a larger function value than B and C and is therefore reflected to D. The new vertex D also has a larger value than B and C and at the next reflection is reflected back to A. The simplex therefore oscillates between ABC and BCD and makes no further progress toward the minimum. The premature "stalling" of the search is a limitation of the simple "reflective-simplex" approach, and the procedure can be improved substantially by using an "adaptive-simplex" that changes its size and shape according to the shape of the local terrain. This has two key advantages as illustrated in Figure 17: (1) the simplex can expand in size when moving through flat uninteresting terrain, which improves the efficiency of optimization, and (2) once it is close to the location of the optimum, it can start to contract, which improves the accuracy of the final solution. For the situation considered in Figures 16 and 17, after just 10 iterations the adaptive simplex is closer to the minimum than the reflective simplex was after 28.

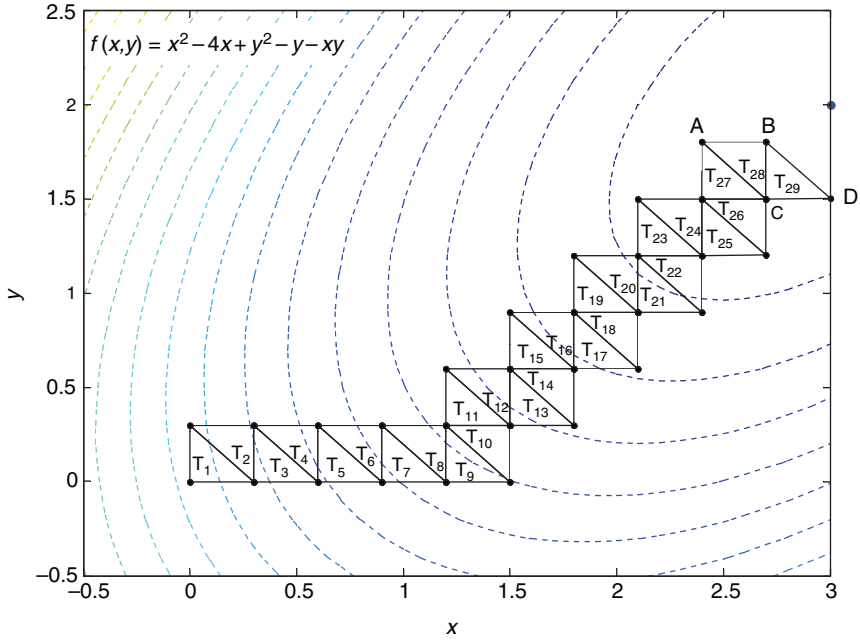

Figure 16 Schematic illustrating the concept of simplex optimization using a simple reflective algorithm, in which the simplex vertex with the largest function value is repeatedly reflected through the line joining the other two vertices. The simplex moves progressively closer to the minimum at (3,2) until the 27th iteration. At this stage, vertex A has a higher function value than vertices B and C, and A is consequently reflected through to D. D also has a larger function value than B and C and so, at the next iteration, is reflected back to A. The simplex therefore makes no further progress toward the minimum and oscillates repeated between ABC and BCD.

A variety of rules have been developed to control the movement and adaptation of the simplex, of which the most famous set is due to Nelder and Mead (Olsson and Nelson, 1975). The Nelder–Mead simplex procedure has been successfully used for a wide range of optimization problems and, due to its simple implementation, is amongst the most widely used of all optimization techniques. Importantly for the current application, simplex optimization is a black-box technique since it uses only the comparative values of the function at the vertices of the simplex to advance the position of the simplex, and it therefore requires no knowledge of the underlying mathematical function. It is also well suited to the optimization of expensive functions since as few as one new measurement is needed to advance the simplex one step. In its usual form, simplex optimization is suitable only for unconstrained optimization, but effective constrained versions have also been developed (Parsons et al., 2007;

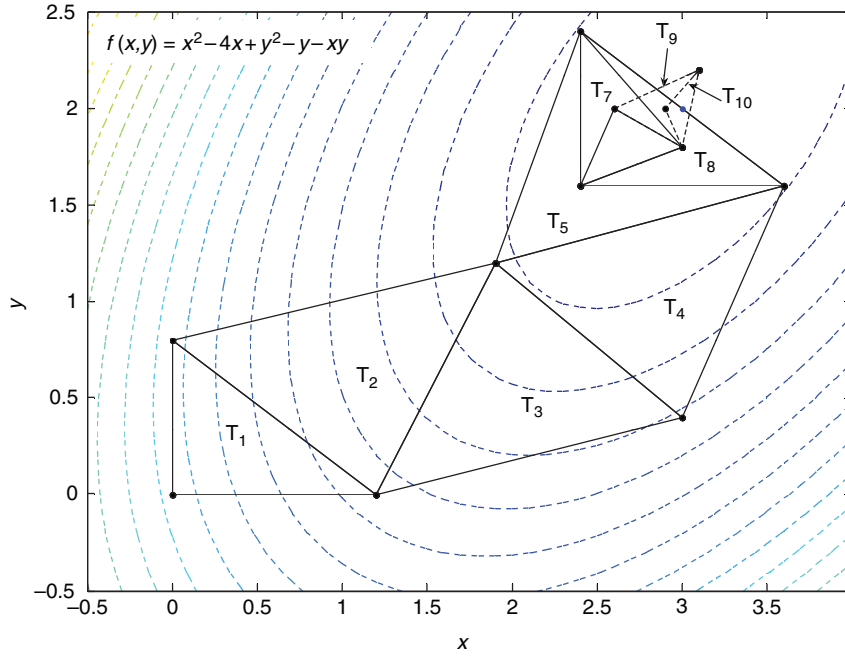

Figure 17 Schematic illustrating the concept of adaptive simplex optimization using the Nelder–Mead algorithm described in Olsson and Nelson (1975). The simplex initially expands in size and so makes rapid progress toward the minimum. It then contracts repeatedly, allowing it to converge on the minimum at (3,2).

Subrahmanyam, 1989). It can therefore be successfully used for constrained optimization of expensive black-box functions. Where it is less effective, however, is in the treatment of noisy functions. If the underlying noise-free function has similar values at the vertices, then the effect of the added noise may be sufficient to change the apparent ordering of the vertices, causing the optimization routine to make a bad decision such as unnecessary contraction (a process that should only happen when the simplex is close to the true minimum). The effect of such a contraction is to bring the vertices closer together, which means the values of the noise-free function at the vertices will be even closer and the distorting effect of noise even worse. The routine may therefore make a series of erroneous decisions that cause it to collapse to a point away from the true minimum.

The key weakness of the standard simplex method is that it only uses information about the function values at the vertices of the most recent simplex, and it entirely disregards data obtained at earlier stages in the optimization. The rejection of historical data as worthless is clearly naïve, but until recently, it has been difficult to see how such data could be

straightforwardly incorporated into a simplex scheme (in which all decisions are based solely on the ranking of the vertices in the current simplex). This problem was solved in 2006 by Martinez (2005), who, building on earlier work by Anderson et al. (2000), proposed a statistical means of ranking the vertices. The detailed mathematical foundations need not concern us here, and it will suffice to say that the approach uses a statistical figure of merit—known as an optimality coefficient—whose value decreases monotonically from zero to minus one as the minimum is approached. In the scheme proposed by Martinez, the simplex optimization is implemented in the usual way, subject to the one change that vertices are ranked according to their relative optimality coefficients (instead of their relative function values). Importantly, the optimality coefficient is calculated using both the function values at the vertices and all historical data, and so the algorithm gains a more global view of the terrain. The statistical simplex approach is far less sensitive to the effects of noise and is therefore in turn less susceptible to becoming trapped at fictitious minima.

The application of the statistical simplex to nanoparticle synthesis is illustrated in Figure 18 using a simple one-dimensional example, in which the statistical simplex algorithm aimed to control the emission wavelength by varying the total flow rate ($F_{tot} = F_A + F_B$) of the injected precursors. (The ratio of the injection rates of the Cd and Se precursors was set equal, that is, $F_A = F_B$, to ensure a constant ratio of Cd : Se in the reaction mixture.) The target wavelength was set to 540 nm, and the vertices of the initial simplex were arbitrarily set to 25 and 35 µl min^{-1}. The variation of the figure of merit, the flow rate, and the peak wavelength are shown as a function of measurement number in Figures 18a and b. In the initial stages of the optimization, the algorithm has no historical data on which to base its decisions and the flow rate oscillates wildly as the simplex reflects backward and forward. The algorithm gradually builds up a picture of the terrain, however, and as it does so the flow rate gradually converges to 5.75 µl min^{-1}, causing the figure of merit to reduce in magnitude toward zero and the peak wavelength to "home-in" on the target of 540 nm. The target is reached after approximately forty measurements. The data in Figure 18b reveal an approximate mirror symmetry between the flow rate F and the peak wavelength λ_{max}. This arises because slower flow rates imply longer reaction times which result in larger particles with longer peak emission wavelengths (and vice versa). In Figure 18c, we show the emission spectra corresponding to the first and last measurements, together with an interim spectrum obtained at measurement ten. The convergence on the target wavelength is clear. (The increase in intensity with measurement number is due to the increasing size of the particles as the emission moves to longer wavelengths, which reduces the quenching effects of surface defects.)

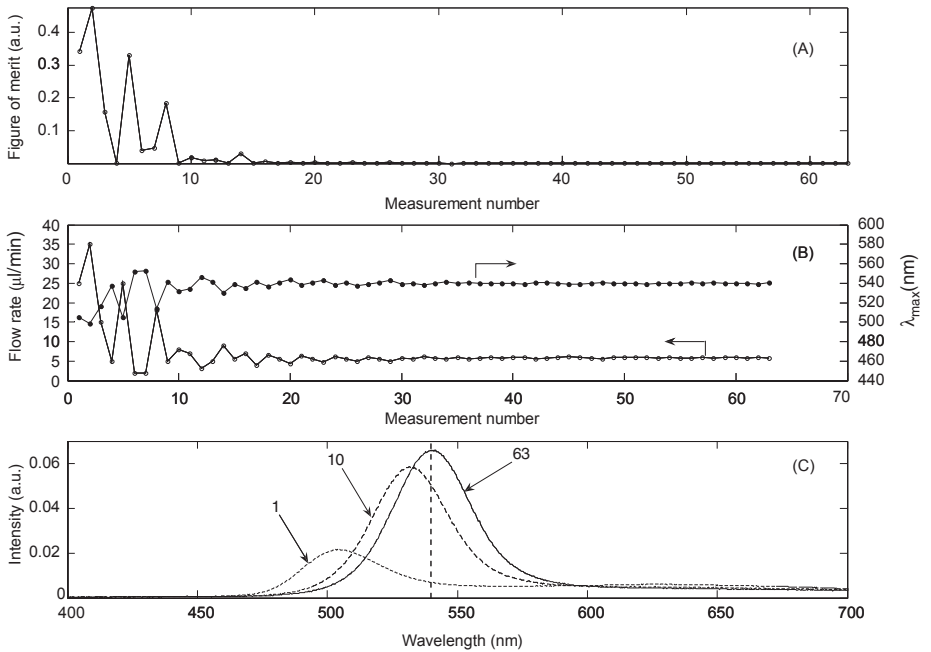

Figure 18 Application of the statistical simplex approach to the one-dimensional optimization of peak emission wavelength, using total flow rate as the sole reaction variable. (A) Variation of the figure of merit with measurement number. (B) Variation of the flow rate and peak wavelength with measurement number. (C) Emission spectra at various stages in the optimization, that is, the initial, tenth, and final measurements. The peak emission wavelength moves progressively closer to the target of 540 nm as the search proceeds.

The above example is intended mainly for the purposes of illustration, and it is arguable that one could just as easily identify the optimum flow rate by simply tweaking the flow rate manually. The automated approach is more useful for higher dimensional problems where the product can be controlled by multiple reaction conditions, for example, by varying the two injection rates of the precursors together with the temperature. In this case, multiple combinations of reaction conditions may yield the same emission wavelength since a faster flow rate (which reduces the reaction time) can, for example, be compensated by a higher reaction temperature (which increases the reaction rate). However, although multiple reaction conditions may yield particles with the same emission wavelength, the emission spectra are liable to have different line widths and intensities. In general one prefers nanoparticles that exhibit sharp intense emission, and it is interesting to ask whether one could optimize the peak wavelength subject to an additional requirement that the particles should, say, emit as brightly as possible. The optimization of multiple properties is known variously as multi-attribute optimization, multi-objective optimization, and decision theory, and is a subject of intense practical interest. There are various ways in which multi-attribute optimization can be carried out, but we will consider here the use of multi-attribute utility functions (MAUFs), that is, utility functions that combine multiple attribute values into a single figure of merit that can then be optimized in the usual way. MAUFs are most straightforwardly created by combining a number of single attribute utility functions (SAUFs) using, for example, a weighted sum. Again, we will not dwell on the mathematical details but simply note that the MAUFs are typically defined to have a value of unity when the multiple attributes all have their worst possible values which reduce to zero if and when they all have their best possible values.

The simultaneous optimization of peak wavelength and intensity is illustrated in Figure 19 for a target wavelength of 550 nm, using the two precursor injection rates as reaction variables. The variation of the figure of merit, the peak wavelength, and the peak intensity with measurement number is shown in Figures 19a–c, respectively. The reaction conditions sampled by the simplex algorithm are shown in Figure 19d, where the vertices of the starting simplex are marked out by a triangle and the final reaction condition is denoted by a grey circle at (3 ml/min, 1 ml/min). The interior region defined by the black dotted lines denotes the constrained parameter space. Figure 19a shows a steady decrease in the value of the figure of merit from an initial value of 0.37 to a final value of 0.2 due to an improvement in both the peak wavelength and the intensity. It is interesting to note that a close match to the wavelength is first reached at measurement 12, but continued searching leads to particles that emit with substantially increased intensity. The emission spectra obtained at measurements 12, 50, and 83 are compared in

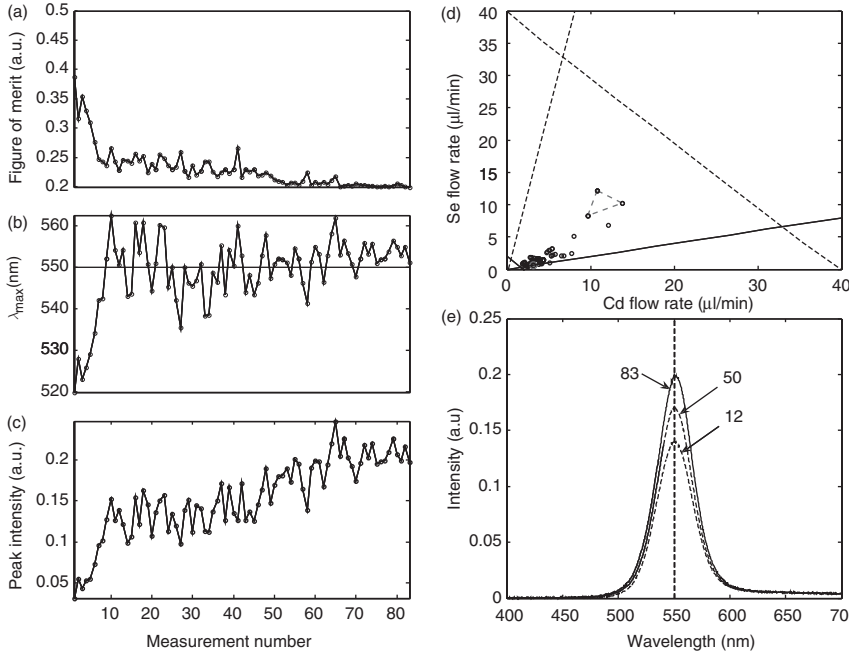

Figure 19 Application of the statistical simplex approach to the simultaneous optimization of peak emission wavelength and intensity, using the individual injection rates of the CdO and Se precursor solutions as the reaction variables. (A) Variation of the figure of merit with measurement number. (B) Variation of the peak wavelength with measurement number. (C) Variation of the intensity with measurement number. (D) Plot showing the conditions sampled by the statistical simplex routine. The initial simplex is denoted by the triangle and the final reaction conditions by the grey circle at (3 ml/min, 1 ml/min). The trapezoidal region defined by the four overlapping dotted lines defines the constrained reaction space. The simplex was only permitted to sample reaction conditions inside this space. (E) Emission spectra at various stages in the optimization, that is, at the 12th, 50th, and 83rd measurement. The peak wavelength shows a close match to the target wavelength in each case but the intensity increases progressively as the search continues, indicating an improvement in particle quality.

Figure 19d, and it is evident that, despite their similar peak wavelengths, the final particles are evidently the most intense.

The above results confirm the feasibility of automating the process of nanoparticle synthesis using simplex-based optimization routines. The approach we outline provides a promising starting point for developing more sophisticated systems that can be used to control a whole variety of nanoparticle properties, including shape, chemical composition, crystal structure, optical properties, and chemical reactivity. In passing, we note that—since the simplex converges on the optimum though a series of

reflections in the general direction of improvement—there is no guarantee it will find the global optimum, which may be located in a distant unexplored region of the parameter space. To stand a better chance of finding the global optimum, different optimization routines are required that work by sampling at discrete locations throughout the parameter space. The search process is divided into two phases: (1) local searching (i.e., solution refinement) in the vicinity of identified optima and (2) global searching in hitherto unexplored regions of the parameter space where superior optima might potentially exist. The interested reader is referred to Krishnadasan et al. (2007) for a discussion of how global optimizers can be applied to nanoparticle synthesis.

5. PROCESS CONTROL

There are two key steps to any practical optimization procedure: firstly, as described in the previous section, the optimal reaction conditions should be identified as efficiently as possible to enable swift production of the desired product, and secondly, having found the (initially) optimal conditions, the reaction conditions should then be regularly updated to compensate for any changes (drift) in the system due, for example, to degradation of the precursor materials or furring of the channel walls. The effects of drift are essentially due to progressive changes in hidden variables over which the user has no control. To compensate for these changes, the active variables must be adjusted accordingly. Hence if, for example, aggressive furring of the channels were to occur, this might significantly reduce the cross-sectional area of the channels and so reduce the residence time for a given flow rate. This could be mitigated either by reducing the total flow rate (to increase the residence time) or increasing the temperature (to increase the reaction rate) or by making appropriate changes to both.

In most cases, the detailed effects of system drift are unpredictable, and it is again important to use black-box algorithms that are able to make adjustments to the reaction conditions on the basis of measured data only without the need for any underlying process model. In order for such a black-box algorithm to determine how to update the reaction conditions, it must constantly dither the reaction conditions about (what it perceives to be) the current optimum. In other words, the algorithm must constantly switch between similar, but slightly different, reaction conditions that yield end products of marginally varying quality. The algorithm can then compare the quality of the particles so obtained and use this information to update the reaction conditions in an appropriate manner. The constant perturbing of the reaction conditions is a necessary activity for tracking the system drift but unavoidably compromises the instantaneous quality of the product since the system frequently finds itself operating at (slightly) non-optimal conditions.

To be effective, the perturbations should be sufficiently small that the product is always of an acceptable quality yet sufficiently large that meaningful comparisons can be made between the different conditions.

Xiong and Jutan have recently reported a real-time algorithm for tracking system drift using a simple adaptation of the reflective simplex approach outlined above (Xiong and Jutan, 2003). In essence, the simplex undergoes repeated reflections that allow it to track the moving optimum. If the optimum is static, the simplex reflects backward and forward and hence remains more or less anchored to the optimum. If the optimum starts to move, however, the repeated reflections drive the simplex in the direction of the drifting optimum.

As noted above, the effects of system drift may be attributed to hidden variables over which the algorithm has no control. In reality, for CdSe nanoparticles synthesized in the manner described above, system drift is a very slow process that occurs over a timescale of days. In order to accelerate the effects of drift and to pose a challenging test for our algorithm, we engineered an artificial situation in which the reaction temperature was programmed to drift linearly with time from 180 to 230 °C over a period of 6 h. (In ordinary circumstances, this would be expected to induce a significant red-shift in the peak emission wavelength.) The system was set to an initial total flow rate of 12 µl min^{-1}, corresponding to particles that initially emitted at 510 nm. The evolution of spectra under fixed and dynamically updated flow-rate conditions was then compared. The upper line in Figure 20 indicates the progressive change in peak wavelength observed over the 6-h period when the flow rate was fixed at 12 µl min^{-1}. The peak wavelength increases almost linearly with time, indicating the formation of larger particles with red-shifted emission (due to the increased temperature-dependent reaction rate, c.f. Figure 10). The peak wavelength shifts by almost 30 nm by the end of the 6-h period. The lower line indicates the variation in the peak wavelength when the flow rate was dynamically updated by the drift-compensating algorithm. The peak wavelength shows an increased volatility due to the constant dithering of the reaction conditions and exhibits a few undesirable outliers in the early stages of the optimization (see circled data points). However, the baseline is effectively static, which indicates that the algorithm is successfully tracking the drifting optimum, and after 6 h the peak wavelength is still approximately 510 nm. The evolution of the flow rate is shown in Figure 21. The flow rate tracks (albeit in a somewhat hap-hazard fashion) the temperature, ensuring that the reaction time decreases sufficiently to offset the increase in temperature. The evolution of the emission spectra with time are shown in Figure 19a and b. The non-stabilized system gives rise to a progressive red shift in the particle emission with time, accompanied by a progressive reduction in the intensity of the defect emission (due to the larger size of the particles). The dynamically stabilized system

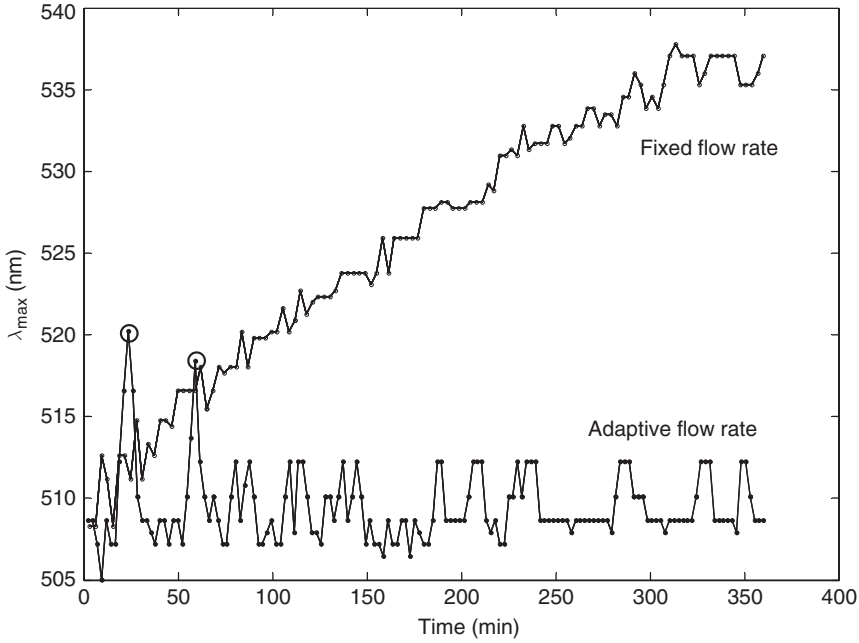

Figure 20 Application of the dynamic simplex to the compensation of system-drift. An artificial example is considered here in which the temperature is ramped linearly with time and the simplex aims to compensate for the changes in the reaction temperature by modifying the flow rate accordingly. The plot compares the change in the peak wavelength when the flow rate is held fixed at its initial value of 12 μl min^{-1} and when it is adapted dynamically by the simplex algorithm. In the former case, the peak wavelength increases steadily with time due to the increasing temperature which increases the growth rate of the particles. In the latter case, the peak wavelength remains fairly close to its initial value of 508 nm.

however shows a much smaller variation in the emission spectra. The remarkable effectiveness of the dynamic stabilization is particularly evident from a comparison of Figure 22c and d, which compare the emission spectra at time 0 and after 6 h for the two cases. The dynamically stabilized system shows a remarkably small change in the emission spectrum, indicating the efficacy of the simplex technique.

6. THE APPLICATION OF AUTOMATED MICROREACTORS IN NANOTOXICOLOGY

In the above sections, we have outlined our recent work in the area of nanoparticle automation and have demonstrated how simple algorithms can be developed to address the two key challenges in nanoparticle

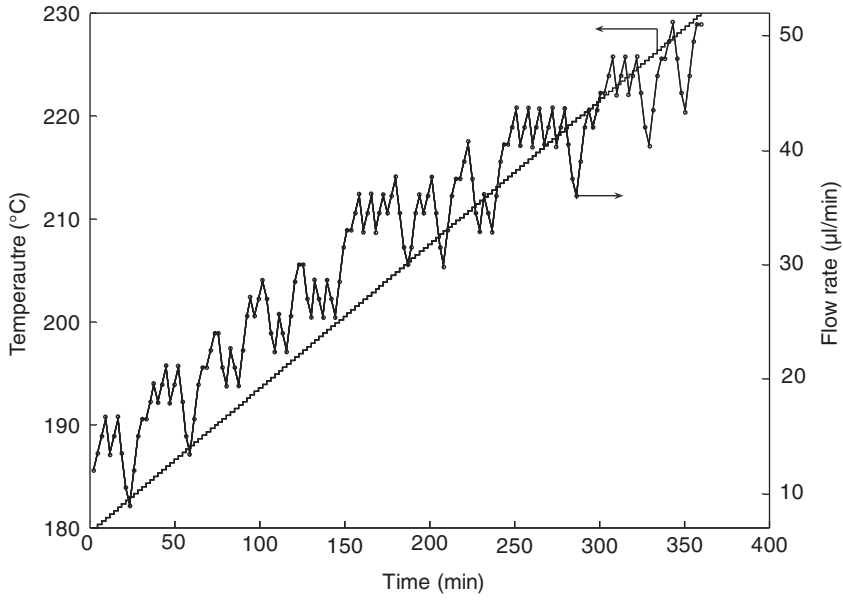

Figure 21 Plot showing how the temperature and flow rate vary with time for the situation considered in Figure 20. The flow rate tracks (albeit in a somewhat hap-hazard fashion) the temperature, ensuring that the reaction time decreases as the temperature increases and that the particles therefore continue to emit with the same approximate peak wavelength.

automation—*fast identification of the optimal reaction conditions and compensation for system-drift*. In this final section, we will spend a few paragraphs outlining how such automated systems might usefully be applied in the specific area of nanotoxicology. In the case of conventional macroscale materials, toxicity is determined primarily by chemical composition. The particle dimensions do not strongly influence the physicochemical properties of the materials themselves, and hence they tend to influence toxicity only indirectly by in essence determining how easily a given particle can reach a susceptible location in the body. In the case of asbestos, for example, the diameter of the fibers determines their ability to penetrate the lung and reach the alveoli where they provoke an inflammatory response that causes deposition of fibrous tissue (and in so doing reduces lung capacity). The thin linear "amphibole" fibers pictured in Figure 23a are able to penetrate the lungs more readily than the curved "serpentine" fibers pictured in Figure 23b and so typically cause more damage. In the case of nanoscale materials, however, the situation is more complex because the particle dimensions also have a strong influence on the physicochemical properties of the particles. In principle, even small variations

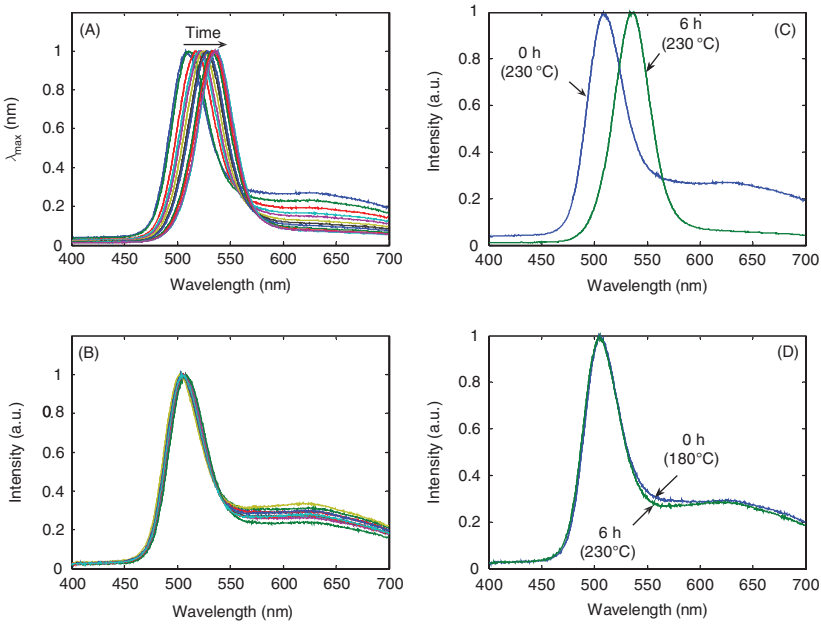

Figure 22 Plots showing how the emission spectra vary with time for the situation considered in Figure 20. (A) The evolution of the spectra for a fixed flow rate of 12 μl min^{-1}. (B) The evolution of the spectra when the flow rates are dynamically updated by the simplex algorithm. (C) Comparison of the initial and final spectra for a fixed flow rate of 12 μl min^{-1}. (D) Comparison of the initial and final spectra when the flow rates are dynamically updated by the simplex algorithm.

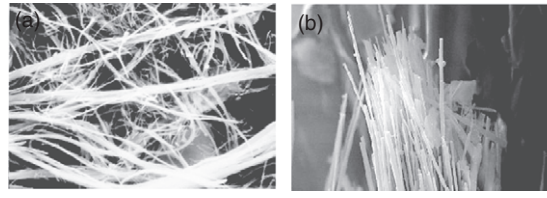

Figure 23 Two varieties of asbestos fibers: (A) amphibole and (B) serpentine. The thin amphibole fibers present a larger toxicological threat due to their greater ability to penetrate lung tissue and reach the alveoli.

in the size and shape of the particles may lead to large variations in their properties and, by implication, their toxicity. Hence, in order to develop a detailed understanding of the toxic effects of a particular nanomaterial, it is important to study its toxicological behavior over the full range of sizes and shapes that are liable to be encountered in practical situations (rather

than a few discrete cases as is usually sufficient for toxicological studies of macroscale particles). Furthermore, a number of other factors are likely to influence toxicity, including crystallinity, surface functionalization, surface charge (zeta potential), hydrophilicity, and propensity for aggregation or agglomeration. The diverse range of factors that can influence the toxicity of nanomaterials presents a formidable challenge for toxicologists, and there is a growing need for new tools and techniques that can be used to analyze the toxicity of nanomaterials.

Conventional macroscale toxicology makes heavy use of libraries of reference materials (RMs)—that is, standard materials with stable well-defined properties that are prepared and handled according to precise protocols. RMs may be investigated by multiple researchers in the confidence that they are all looking at essentially the same materials, enabling meaningful comparisons to be made between their research findings. There is now an urgent need to develop similar libraries for nanotoxicology, in which commonly used nanomaterials are provided in as wide a variety of sizes, shapes, and surface functionalizations as possible. A number of recent studies have investigated the feasibility of developing such libraries, and there is an emerging consensus about the requirements of such libraries. Firstly, RMs should be available for all the most commonly used nanomaterials, namely, metals, metal alloys, metal oxides, quantum dots, and fullerenes. Secondly, the RMs should be provided in a range of monodisperse sizes and shapes. Thirdly, the RMs should be of verifiable quality with minimal batch-to-batch variation between nominally identical materials. And finally, the RMs should be available in kilogram quantities to ensure sufficient material is available for worldwide studies. In addition, the studies drew attention to one issue of particular concern, namely, the significant differences in reactivity (and hence potential toxicity) that frequently exist between freshly prepared and stored particles. This issue throws into serious doubt the viability of the conventional RM approach—in which standardized materials are prepared in centralized facilities and then transported to their intended location of use—since the particles are liable to "age" substantially in transit. Ideally, the nanomaterials should be prepared on demand at the point of use, but this is inconvenient, potentially dangerous, and requires a level of synthetic expertise that the typical practicing toxicologist does not have. The method outlined in this chapter, however, offers a possible solution. The automated systems described above could be readily integrated into a simple self-contained desktop machine that produces nanomaterials on demand to a well-defined and verifiable specification. In this way, nanomaterials could be prepared as and when required, enabling reliable toxicological studies to be undertaken on both fresh and aged samples. Moreover, the high levels of process control afforded by the microfluidic format enable systematic studies in which key nanomaterials

properties—size, shape, chemical composition, zeta potential, etc.—may be varied with ease. At the time of writing, a wide variety of nanomaterials have been successfully prepared in microreactors, including CdSe, CdS TiO2, Ag, Au, and Co to name but a few. In fact, of the key families of nanomaterials identified above, only fullerenes have yet to be successfully synthesized in a microfluidic format (due to the normal use of combustion-based preparation routes). The automated microfluidic route is therefore an attractive means of making most of the commonly used families of nanomaterials. Importantly, it satisfies most of the other requirements of RMs noted above: the high levels of process control enable the preparation of size-monodisperse series of nanomaterials and the use of online monitoring allows quality assurance and minimizes batch-to-batch variation.

The major objection to using the microfluidic approach might seem to be inadequate materials throughput. In fact, despite the small size of the reaction channels, surprisingly large quantities of materials can be prepared using microfluidic devices. In the case of CdSe quantum dots, for example, the nanoparticles can be prepared using flow rates of several hundred microliters per minute which, for typical precursor concentrations of $10\,\mu g\,\mu l^{-1}$ implies production rates of order $1\,mg\,min^{-1}$. Moreover, due to the small size of the channels, it is possible to integrate multiple channels onto a single microchip, enabling the syntheses to be carried out in parallel. This approach is called "scale-out" in contrast to the usual process of "scale-up" and has the important advantage that all channels can be fabricated with identical dimensions, and the increased throughput can therefore be achieved without the need to reformulate the chemical reactions in any way (as is invariably necessary for scale-up). In a small multichannel system containing just twenty parallel reactors, production rates of about $1\,gh^{-1}$ are easily achievable which is sufficient for most practical applications. Hence, despite the small size of the microreactors, an appreciable amount of material can be produced in a relatively short period of time.

In our current work, we are aiming to develop parallel systems of this nature for a variety of applications where on-demand production of high-quality nanoparticles is required.

7. CONCLUSIONS

It is fair to say that over the last 8 years there has been a significant growth in microfluidic-based methods for synthesizing nanoscale materials. Significantly, such materials are of a quality that matches or exceeds materials produced using traditional bulk methods. Furthermore, the continued development of integrated multicomponent systems will create invaluable tools for improving both the properties and yields of this important class of materials.

More specifically, the described approaches for reaction automation represent an important first step toward both simplifying and automating the process of nanoparticle production, but there is still considerable scope for improvement. To date, we have applied our approaches to a relatively small number of nanoparticle syntheses using very simple chemical routes. However, some extremely sophisticated chemical syntheses have been developed in recent years that are capable of producing near defect-free nanoparticles of extremely high monodispersity, and it will be important to adapt and apply our automation procedures to these synthesis routes also. In addition, although in this chapter we have focused on the synthesis of fluorescent quantum dots, the general strategy we describe has wider applicability to any particles whose properties can be monitored (directly or indirectly) inline. In the case of nonfluorescent dots and rods, dynamic light scattering is likely to provide an especially effective tool for controlling the size, shapes, and dispersity of the particles. As mentioned above, careful design of the control algorithms is crucial to achieving effective automation, and considerable scope exists for achieving further efficiency and control through improved algorithm design. Notwithstanding these challenges, we consider the general approach outlined above to offer a powerful route to the automated production of optimized nanoparticles which has the potential to transform the efficacy of nanoparticle synthesis in terms of control, yield, and ease-of-use.

ACKNOWLEDGEMENT

We acknowledge financial support from the European Commission (NanoReTox, http://www.nanoretox.eu) and the Royal Society under its Industry Fellowship scheme.

REFERENCES

Alivisatos, A. P. *Science* **271**, 933–937 (1996).
Anderson, B., Moore, A., and Cohn, D. A. "Presented at International Conference on Machine Learning 2000", Stanford University, CA, USA (2000).
Boleininger, J., Kurz, A., Reuss, V., and Sonnichsen, C. *Phys. Chem. Chem. Phys.* **8**, 3824–3827 (2006).
Chan, E. M., Alivisatos, A. P., and Mathies, R. A. *J. Am. Chem. Soc.* **127**, 13854–13861 (2005).
Chan, E. M., Mathies, R. A., and Alivisatos, A. P. *Nano Lett.* **3**, 199–201 (2003).
Cottam, B. F., Krishnadasan, S., deMello, A. J., deMello, J. C., and Shaffer, M. S.P. *Lab. Chip* **7**, 167–169 (2007).
deMello, A. J. *Nature* **442**, 394–402 (2006).
DeMello, J., and DeMello, A. *Lab Chip* **4**, 11N–15N (2004).
Donega, C. D., Hickey, S. G., Wuister, S. F., Vanmaekelbergh, D., and Meijerink, A. *J. Phys. Chem. B* **107**, 489–496 (2003).
Dumais, P., Callender, C. L., Ledderhof, C. J., and Noad, J. P. *Appl. Opt.* **45**, 9182–9190 (2006).

Dushkin, C. D., Saita, S., Yoshie, K., and Yamaguchi, Y. *Adv. Colloid Interface Sci.* **88**, 37–78 (2000).
Edel, J. B., Fortt, R., deMello, J. C., and deMello, A. J. *Chem. Commun.* 1136–1137 (2002).
Euliss, L. E., DuPont, J. A., Gratton, S., and DeSimone, J. *Chem. Soc. Rev.* **35**, 1095–1104 (2006).
Green, M. *Angew. Chem. Int. Ed.* **43**, 4129–4131 (2004).
Han, G., Ghosh, P., and Rotello, V. M. *Nanomedicine* **2**, 113–123 (2007).
Hung, L.-H., Choi, K. M., Tseng, W.-Y., Tan, Y.-C., Shea, K. J., and Lee, A. P. *Lab Chip* **6**, 174–178 (2006).
Iliescu, C. *Informacije Midem-J. Microelectron. Electron. Compon. Mater.* **36**, 204–211 (2006).
Kawazoe, T., Yatsui, T., and Ohtsu, M. *J. Non-Crystall. Solids* **352**, 2492–2495 (2006).
Khan, S. A., Gunther, A., Schmidt, M. A., and Jensen, K. F. *Langmuir* **20**, 8604–8611 (2004).
Klostranec, J. M., and Chan, W. C. W. *Adv. Mater.* **18**, 1953–1964 (2006).
Kohler, J. M., Wagner, J., and Albert, J. *J. Mater. Chem.* **15**, 1924–1930 (2005).
Kolda, T. G., Lewis, R. M., and Torczon, V. *Siam Rev.* **45**, 385–482 (2003).
Koo, J. M., and Kleinstreuer, C. *J. Micromech. Microeng.* **13**, 568–579 (2003).
Krishnadasan, S., Brown, R., deMello, A. J., and deMello, J. C. *Lab Chip* **7**, 1434–1441 (2007).
Krishnadasan, S., Tovilla, J., Vilar, R., deMello, A. J., and deMello, J. C. *J. Mater. Chem.* **14**, 2655–2660 (2004).
La Mer, V. K., and Dinegar, R. H. *J. Am. Chem. Soc.* **72**, 4847–4854 (1950).
Malik, M. A., O'Brien, P., and Revaprasadu, N. *Phosphorus Sulfur and Silicon and the Relat. Elem.* **180**, 689–712 (2005).
Manz, A., Harrison, D. J., Verpoorte, E. M. J., Fettinger, J. C., Paulus, A., Ludi, H., and Widmer, H. M. *J. Chromatogr.* **593**, 253–258 (1992).
Martinez, E. C. *Indus. Eng. Chem. Res.* **44**, 8796–8805 (2005).
Masala, O., and Seshadri, R. *Annu. Rev. Mater. Res.* **34**, 41–81 (2004).
Matsui, I. *J. Chem. Eng. Jpn* **38**, 535–546 (2005).
Medina, C., Santos-Martinez, M. J., Radomski, A., Corrigan, O. I., and Radomski, M. W. *Br. J. Pharmacol.* **150**, 552–558 (2007).
Milliron, D. J., Hughes, S., and Alivisatos, A. P. *Abstr. Pap. Am. Chem. Soc.* **227**, U274–U274 (2004).
Millman, J. R., Bhatt, K. H., Prevo, B. G., and Velev, O. D. *Nat. Mater.* **4**, 98–102 (2005).
Murray, C. B., Norris, D. J., and Bawendi, M. G. *J. Am. Chem. Soc.* **115**, 8706–8715 (1993).
Nakamura, H., Yamaguchi, Y., Miyazaki, M., Maeda, H., Uehara, M., and Mulvaney, P. *Chem. Commun.* 2844–2845 (2002).
Nightingale, A. M., and de Mello, J. C. *Chem. Phys. Chem.* **10**, 2612–2614 (2009).
Olsson, D. M., and Nelson, L. S. *Technometrics* **17**, 45–51 (1975).
Ozin, G., and Arsenault, A. *Nanochem. A Chem. Approach Nanomater.* (2005).
Park, J., An, K. J., Hwang, Y. S., Park, J. G., Noh, H. J., Kim, J. Y., Park, J. H., Hwang, N. M., and Hyeon, T. *Nat. Mater.* **3**, 891–895 (2004).
Parsons, D. J., Green, D. M., Schofield, C. P., and Whittemore, C. T. *Biosyst. Eng.* **96**, 257–266 (2007).
Peng, Z. A., and Peng, X. G. *J. Am. Chem. Soc.* 123, 183–184 (2001).
Pradhan, N., Katz, B., and Efrima, S. *J. Phys. Chem. B* **107**, 13843–13854 (2003).
Rao, C. N. R., Agrawal, V. V., Biswas, K., Gautam, U. K., Ghosh, M., Govindaraj, A., Kulkarni, G. U., Kalyanikutty, K. R., Sardar, K., and Vivekchandi, S. R. C. *Pure Appl. Chem.* **78**, 1619–1650 (2006).
Rhyner, M. N., Smith, A. M., Gao, X. H., Mao, H., Yang, L. L., and Nie, S. M. *Nanomedicine* **1**, 209–217 (2006).
Shalom, D., Wootton, R. C. R., Winkle, R. F., Cottam, B. F., Vilar, R., deMello, A. J., and Wilde, C. P. *Mater. Lett.* **61**, 1146–1150 (2007).
Shestopalov, I., Tice, J. D., and Ismagilov, R. F. *Lab Chip* **4**, 316–321 (2004).

Song, Y. J., Modrow, H., Henry, L. L., Saw, C. K., Doomes, E. E., Palshin, V., Hormes, J., and Kumar, C. *Chem. Mater.* **18**, 2817–2827 (2006).
Subrahmanyam, M. B. *J. Optim. Theory Appl.* **62**, 311–319 (1989).
Takagi, M., Maki, T., Miyahara, M., and Mae, K. *Chem. Eng. J.* **101**, 269–276 (2004).
Uehara, M., Nakamura, H., and Maeda, H. *J. Nanosci. Nanotechnol.* **9**, 577–583 (2009).
Vilkner, T., Janasek, D., and Manz, A. *Anal. Chem.* **76**, 3373–3385 (2004).
Wagner, J., and Kohler, J. M. *Nano Lett.* **5**, 685–691 (2005).
Wang, H. Z., Li, X. Y., Uehara, M., Yamaguchi, Y., Nakamura, H., Miyazaki, M. P., Shimizu, H., and Maeda, H. *Chem. Commun.* 48–49 (2004).
Wang, H. Z., Nakamura, H., Uehara, M., Yamaguchi, Y., Miyazaki, M., and Maeda, H. *Adv. Funct. Mater.* **15**, 603–608 (2005).
Xia, D. Y., Li, D., Ku, Z. Y., Luo, Y., and Brueck, S. R.J. *Langmuir* **23**, 5377–5385 (2007).
Xiong, Q., and Jutan, A. *Chem. Eng. Sci.* **58**, 3817–3828 (2003).
Xue, Z. L., Terepka, A. D., and Hong, Y. *Nano Lett.* **4**, 2227–2232 (2004).
Yang, H., Luan, W., Wan, Z., Tu, S-t., Yuan, W.-K., and Wang, Z. M. *Cryst. Growth Des.* **9**, 4807–4813 (2009).
Yen, B. K. H., Günther, A., Schmidt, M. A., Jensen, K. F., and Bawendi, M. G. A. *Angew. Chem. Int. Ed.* **44**, 5447–5451 (2005).
Yen, B. K. H., Stott, N. E., Jensen, K. F., and Bawendi, M. G. *Adv. Mater.* **15**, 1858–1862 (2003).
Yordanov, G. G., Gicheva, G. D., Bochev, B. H., Dushkin, C. D., and Adachi, E. *Coll. Surf. A-Physicochem. Eng. Asp.* **273**, 10–15 (2006).