A 'WHOLE SYSTEM DESIGN' APPROACH TO PERSONAL TRANSPORT

Hugo Spowers

OSCar Automotive Ltd., UK.
hugospow@compuserve.com

"It is not necessary to change. Survival is not mandatory."

W. Edwards Deming

Introduction

This paper summarises OSCar's proposals for the development of sustainable personal transport vehicles, key elements of a strategy for their delivery and some pragmatic observations, principally on minimising the risk that a less sustainable ubiquitous architecture should get there first.

The principal contention is that, contrary to received wisdom, fuel cell technology is sufficiently advanced to do what needs to be done; the popularly perceived technological barriers are artifices of the approach necessarily adopted by the automotive industry, whereas the absolute barriers are due to our systems of politics, business and thinking.

The vehicle technology itself is covered first, as the obvious starting point and the easiest problem to solve. This is followed by a discussion of the disruptive nature of the transition to a fuel cell powered automotive world. Going up a level of design, to redesigning processes, the harder problem of business strategies to deliver this solution is discussed, again derived through Whole System Design driven by the characteristics of the vehicle technology. Rather than cover all relevant strategies thinly, the paper concentrates on the design process and the ownership model; manufacturing strategy, distribution and governance have not been addressed here. Exhaustive discussion of the whole universe of competing possibilities has also been omitted, but notes have been included at the end on the most commonly voiced questions.

Purpose, principles and criteria

The Purpose behind these proposals is to develop a solution to the requirement for sustainable personal transport, to include not just the vehicle technology but also the sustainable delivery of that service and resilient organisations to provide it. By way of clarification of what has driven the development of these proposals, various drivers, principles and criteria are outlined below that have informed the process.

Drivers

- The above purpose is defined in full acknowledgement that the balance between personal and mass transport must change, but that there will always be a personal transport element in a sustainable system.
- Sustainable is taken to mean no net loss of capital, especially natural, so we must meet transport demand on the income available (renewable resources).
- The primary (not only) aspect is energy efficiency in use; energy efficiency will become orders of magnitude more important than ever it has been in our capital based transport system.

- The problem is urgent so we cannot rely on distant technologies, yet there is no point in investing in solutions that have no hope of ever achieving the goal; being 'less bad' is no good.

Design principles and criteria

- Resilience is more important than profit; by definition, no organisation can be resilient without profit anyway, but it can be profitable without resilience.
- There is an emphasis on 'aikido' engineering and strategies that maximise opportunities and minimise barriers and resource requirements, particularly financial.
- Self-regulating systems are more efficient and more effective (but more difficult to design) than 'policed' systems.
- Whole system design (evaluation of alternatives at the system level, not that of the subsystem) applies to all aspects of the design of these proposals - the car, transport system, value network that delivers it, corporate structure and governance of organisations involved, legal relationships binding them together into an industry and a strategy for getting there from here.
- In order to ensure that all investment contributes to the goal and to avoid multiple disruptions, solutions have been developed by backcasting from the end goal defined above.
- Any plan must be designed to accommodate the reality that if successful, it will not happen overnight, and also flexible enough to accommodate changing conditions and the emergence of new possibilities - YET we do have to make some choices NOW.

Efficient hydrogen fuel cell vehicle technology

"But there are times you need the courage to take a great leap;
you can't cross a chasm in two small jumps."

David Lloyd George

OSCar is at present running two hydrogen fuel cell car programmes, both vehicles to be running in June/July 2007. The first is a Morgan sports car, the LIFECar [1], part-funded by the DTI, and the second is an urban car, Hyrban [2]. Cranfield and Oxford Universities are partners in both projects and the LIFECar consortium also includes the Morgan Motor Co., Qinetiq and BOC. These cars are premised on the principal ideas articulated by the Rocky Mountain Institute in the Hypercar programme [3]. In both cases, the objectives are to demonstrate that fuel cell vehicles are a practical proposition using best available off-the-shelf technology and that a step change in efficiency (factor 4 to 6) is possible by designing a car for a fuel cell rather than grafting one into a car designed for an engine.

None of the components or technologies make any performance or economic sense on their own, and none compare favourably with their approximate equivalents in an internal combustion engine (ICE) car. Yet we expect Hyrban to have an urban duty cycle fuel efficiency of 250mpg equivalent, a factor 5+ improvement on the Smart. Two key phenomena, working in symphony, enable this - mass decompounding and the decoupling of steady state and transient loads.

Weight and fuel consumption are tightly coupled. Advanced composite materials offer the most efficient structural solutions. Although expensive, prices are falling, the weight of composites required will be low and they will be amortised over a much longer period - of which more in 'Sale of service.' By reducing the weight of any subsystem, the structural weight reduces, as it is dictated by

[1] http://www.autoindustry.co.uk/news/22-06-05_6
[2] http://www.boc.com/news/2005/article_977_05dec05.asp
[3] Lovins, A.B. and Cramer, D.R., 2004, Hydrogen, hypercars and the automotive transition, Int. J. Vehicle Design, Vol 35, Nos. 1/2, pp 50-85.

the mass of the heavy components whose kinetic energy the structure has to manage in any impact. In a conventional car, the maximum constant demand is when cruising on a motorway, but it is still only about 20% of peak power; the remaining 80% is only required under acceleration. Consequently, for 90% of its working life, the car is carrying 80% redundant engine mass - and excess weight in most other systems. If kinetic energy is recovered under braking, and used to supplement a fuel cell under acceleration, the fuel cell only need be sized for constant demand, structural weight drops further and energy demand drops as well. There are also step changes in weight reduction as the lower weight and narrower tyres eliminate the need for powered systems, such as brakes and steering. This iterative process of mass decompounding enables a vehicle weight reduction of approximately 50%, for no loss of functionality.

The regenerated energy captured by wheel motors under braking is stored in our cars in ultracapacitors, although other alternatives can be used. The models that we have developed at Cranfield University allow for complex load-levelling systems; if designing an urban car for the hills of San Francisco for instance, the network will incorporate lithium ion batteries as well as capacitors as the optimal solution for the duty cycle in that city. The round trip energy efficiency of ultracaps is very high and their power density is much greater than batteries, so they can absorb very high currents under braking. It must be emphasised that the significance of the regenerated energy is not that it is saved and reused later but that it enables the decoupling of demand, the downsizing of the fuel cell and thus weight reduction. However, for this decoupling to accommodate real world duty cycles, the regeneration must be highly efficient; no energy can be thrown away with friction brakes, or else the performance will degrade with each successive acceleration event. Effective braking can only be achieved by four-wheel braking, so the vehicle needs four motors for the principle to work (two motors would introduce packaging constraints and, more importantly, differential issues). The result is a Network Hybrid, in which all powertrain components are connected to each other in a network and energy can flow in either direction on any path, except back into the fuel cell.

Disruptive change in the automotive industry

"Let the mind be enlarged, according to its capacity, to the grandeur of the mysteries, and not the mysteries contracted to the narrowness of the mind."

Sir Francis Bacon

Fuel cells pose a classic opportunity for a disruptive technological shift, as characterised by Clayton Christensen [4] - a step change in a technology that performs worse than the incumbent on the metrics by which it is sold. The reward that the Network Hybrid offers is a vehicle architecture that will be more efficient by a factor of 4 to 6 than the grafting of a fuel cell into a conventional linear powertrain in a steel body. Yet cars are not sold on efficiency but on performance and, increasingly, the luxurious automated nature of the private space.

The vast body of opinion within the auto industry sees fuel cell vehicles as the only viable end game but, compared with the ICE, fuel cell power density and cost are not remotely competitive with the internal combustion engine. There are two solutions to this conundrum:-

- Improve the technology so that it is competitive
    - Success through basic science; large teams, expensive and slow
- Accept the limitations of the technology and design around them
    - Success through design; small teams, cheap and quick

Whilst I am endorsing a focus on the pattern of relationships within the vehicle, there is no doubt that the industry has very little room to manoeuvre due to the marginal profitability characteristic of a

[4] Christensen, Clayton, 1999, The Innovator's Dilemma, HBS.

commodity industry, so has to plot an incremental path forward. Their principal constraint is the sunk investment, financial, strategic and cultural, in pressed steel body construction, with the weird economics that accompanies it. In order to smooth the path from what they do now to such a fuel cell future, the industry has very little option but to graft a fuel cell into the vehicle technology they purvey. To achieve this, they must throw money at the basic science, with three principal objectives:

- Increase power density
- Reduce fuel cell cost
- Increase hydrogen storage density

Unfortunately, the first two objectives conflict badly, as power density is increased by pressurising the stacks, increasing the subsystem count, stresses, costs and losses; this also magnifies the storage problem. Note also that both of these objectives are about reducing volume, not weight; this has a compounding effect, as heavier cars exacerbate the problems. The inescapable observation is that the key barriers to commercialisation, fuel cell cost and hydrogen storage, are artificially generated by this strategy:-

- Fuel cell cost is pro rata to power, but cars could use one fifth of the power
- Existing hydrogen storage is incapable of storing enough hydrogen to give a 300 mile range at 35 miles per kilo

However, there is no doubt that something akin to a Ford Focus will eventually be brought to market with a linear powertrain and a 300 mile range at a viable price - at a high development cost in both time and dollars. Once that happens, it will become very much harder ever again to readdress the basic architecture (which one has to remember has been dominant for over 100 years) and its inherent inefficiency, a major handicap if we are attempting to develop a sustainable transport solution. There is therefore a window of opportunity - say 10 to 15 years - in which not necessarily to achieve ubiquity, but to successfully build a bridgehead into the market.

This window is particularly important and exciting in light of the fact that the Network Hybrid approach is able to deliver a commercially viable product at a much earlier stage of fuel cell development than the traditional linear powertrain. Fuel cell technology is at the beginning of the 'S-curve' of development. Once there is a revenue stream to support it, the technology will improve much more rapidly and will move inexorably up through the market sectors in which the technology cannot yet compete effectively.

Sale of service as central to the design of a fuel cell car industry

"Nothing is less productive than to make more efficient what should not be done at all."

Peter Drucker

The business model proposed is based on leasing rather than selling or, more strictly, sale of service rather than product. The sale of product rewards the maximisation of resource and energy throughput; we cannot hope to develop the technologies and products that a resource-efficient, cyclical industrial society requires by rewarding the opposite of that for which we strive and then relying on regulation as our damage limitation policy. In contrast, the sale of service rewards the minimisation of resource throughput, a good example of a self-regulating system, both more efficient and more effective. However, sale of service must be embedded at a system level in an organisation, rather than applied at a cosmetic level. As currently practised in the automotive industry, leasing is merely an alternative form of contract with a customer, designed to shift more product.

If selling a fully bundled transport service, the manufacturer has an interest in the car being as simple, cheap, low maintenance and reliable as possible, the consumer keeping it as long as possible and it being as secure and fuel efficient as possible (as the company owns the asset and also pays all the fuel bills incurred). All these are the opposite of current manufacturers' interests, so this model aligns the previously opposed interests of consumer and manufacturer.

Conventional cars have an economic life of only 4 years, beyond which manufacturers are not prepared to support them. If one draws an envelope around the provision of a vehicle over 20 years, the total cost of provision of that service is much lower with one vehicle that lasts twenty years than five vehicles that last four years, even if they are more expensive to build and one allows for interest. Such a manufacturer can:-

1. take a greater margin yet offer a more expensive product to the consumer at a lower price
2. capture the same quantity of revenue streams with 80% less production capacity
3. access the revenue streams that the secondhand market offers [5]

Conventional cars corrode and have a large number of mechanical components that rub against each other, need lubricants that need replacing and still wear out. Furthermore, there is little residual value at end of life, as it resides in machining tolerances. This is clearly not a promising candidate for leasing, so selling them makes sense.

It is true that there have been significant advances in recent years in longevity, but ownership cycles have continued to fall. This conjuring trick has been achieved through marketing together with complex financial mechanisms, that generate the illusion of car replacement every two years as the cheapest solution [6], and very expensive electronic units that are sealed, a significant driver for disposal of out-of-warranty cars and ever lower second-hand car values.

On the other hand, a composite-bodied fuel cell car is made of inert materials, so there is no structural degradation other than wear and tear, and there are no moving components in the car other than the electric motors and the secondary braking system. The motors have no rubbing contact and no lubricants, apart from the main bearings which are also the wheel bearings, and the brakes will never require replacement, because the motors do most of the braking and the vehicle is so much lighter anyway. This is a much more promising candidate for leasing than the incumbent technology.

There are some necessary conditions, further benefits and barriers that arise from the sale-of-service model and are closely linked:-

- Retention of asset - Fundamental to this model working is that the manufacturer retains ownership of the asset and its maintenance liability if the self-regulating drivers are to remain intact.
- Residual value - The vehicle will have higher residual value to the manufacturer than anyone else if it is designed to be recycled - as indeed it must be - so it does not make sense to sell it.
- Market penetration - Elimination of the capital cost of the vehicle will aid penetration of the market, especially with new technology that is relatively expensive.
- Working capital - One downside of a product with a built-in long term revenue stream is that the cost of manufacturing precedes the income by a considerable margin. Consequently, there is an enormous working capital requirement, which grows more quickly than the revenue stream in the early growth phase; as a new technology, it is likely to receive an especially lukewarm reception from the financial world.
- Cost of capital - The conventional argument against leasing is that it is always more expensive; IBM's management manuals insist on always making a purchase, as their cost of

[5] The existing industry is acutely aware that these are greater than those of the new car market - and that they are incapable of accessing them.

[6] This is especially apparent when the market softens, as the steel car industry is remarkably inflexible in its output. The margins are tiny, the fixed costs of production are very high, the variable costs very low and there is excess capacity globally; this combination is very inconvenient when demand is variable.

capital is lower than anyone else's. This is only true if the photocopier one chooses to buy is the same one that one could lease; if sale-of-service is embedded upstream as far as design, the product that emerges will be fundamentally different - it would be strategically daft to design such a car if one was then to proceed to sell it.

- End of life liability - The concept of waste, the most profound and novel invention of the industrial age, is merely evidence of a loop yet to be closed. End of life liabilities are one of the European motor industry's three most fundamental concerns, along with margins and block exemption, and the easiest way to reduce them is to reduce the level of production necessary to service a given share of the market.

The Shift from Linear to Cyclical Systems of Production

For industry interests to be truly aligned with our need for resource efficiency, one further adaptation is required - upstream extension of 'sale of service'. Many further benefits accrue, and the barriers are lowered, if the concept is extended into the supplier network that supports the industry. Fuel cells, electric motors, even structural materials and interior trim could be supplied to the manufacturer on the same basis as the car is supplied to the consumer, ie the supplier retains the asset and is paid by regular monthly instalments:-

- Subsystem and component suppliers are influenced in the same way as the car manufacturer, rewarded for longevity, efficiency and reliability; if they were purchased by the manufacturer, the supplier would have the same direct interest in reducing product life.
- End-of-life liabilities can be reduced and handled more effectively, as they are handed back to suppliers who would thus have a direct interest in designing for true recycling.
- As with the car to the consumer, it is easier for subsystem manufacturers to penetrate the automotive sector with expensive technologies, components and materials, and the industry would benefit if they succeeded.
- The major inconvenience for the sale of service model is dramatically reduced - the huge working capital requirement for production of vehicles with a payback in excess of a decade. The working capital requirement is distributed around the value network within which the car manufacturer is embedded; the working capital must still be found of course, but the scale is reduced for each partner in the final product and is net of each partner's margin.

An example of this would be recyclable structural materials, as opposed to downcyclable [7]. Shell or Du Pont could develop and <u>sell</u> totally recyclable composite resins and solvents; displacing steel as the primary structural material of the automotive sector might be a very tempting market for them - but they would be very unwise to do so. Once the global production level of composite cars were to stabilise, the same 'technical nutrients' would be going round in endless loops; rather than soften, the market would collapse. On the other hand, if they were to lease the materials as a service to the car manufacturer, and remained responsible for the final recovery of the body structure chemicals, they could increase turnover whilst simultaneously reducing inputs. This could be particularly attractive for the oil industry - in the light of ever-increasing pressure, from all quarters, to reduce extraction of oil.

The same model also has significant appeal for the emerging fuel cell industry, as their forecast sales have consistently failed to materialise. Fuel cells are universally regarded as too expensive, yet sale of installed kilowatt hours would all but eliminate this entry barrier. Furthermore, unlike a petrol engine whose value is in machining tolerances, the bulk of the value in the cell is recoverable at end of life; platinum as a catalyst is not consumed, at all, and, with no moving parts, the intricate bipolar plates

[7] Most recycling is in fact downcycling; the output cannot be put back in as input and used for the same purpose, as it is of lower quality. Recycling implies total recovery of the input and its attributes.

only need cleaning. If aiming for a sustainable business model, it makes much more sense to lease the cell and retain ownership of the limiting factors.

The same also applies to the electric motors; they have high value, non-consumable components, in this case the permanent magnets. Finally, another major element in the car that could be leased is the interior. Wear and tear may lead to a shorter life for the interior than the powertrain or the structure, which will affect the desirability of a ten-year old car released back into the market. If the interior were to be made of recyclable materials, the supplier can be responsible for the reclamation of materials and a prospective new owner, for a modest premium, can have a new interior in his chosen colour scheme. These all appear to be more attractive business propositions for suppliers.

However great society's need for such technologies, components and materials, we cannot expect the private sector to develop them as long as the product continues to be sold [8]. However, unlike many changes in business practice, we don't have to wait for the rest of industry, or even direct competitors, to adopt this model, as doing so unilaterally turns a potential problem into a source of competitive advantage immediately.

Open Source Design

> "All truth passes through three stages.
> First, it is ridiculed.
> Second, it is violently opposed.
> Third, it is accepted as being self-evident."
>
> Arthur Schopenhauer

OSCar is currently engaged in developing an Open Source (OS) approach to the development of these vehicles. There are a number of generic arguments for OS being an attractive alternative to IP as a foundation for the conventional corporate mission of return on capital. These arguments are summarised below but an argument is then articulated as to why OS presents the optimal chance, in the specific circumstances surrounding fuel cell technology in the automotive sector today, of the impending disruption successfully leading to a sustainable future.

The pursuit and enforcement of IPR is demanding of resources, financial and human, whilst the cost of development through open source is extremely low. Open source depends on largely voluntary participation, or the participation of organisations with a specialised (not necessarily altruistic) interest in achieving specific goals. The costs are therefore spread, and are far less than those that would be spent on an equivalent closed enterprise. Input necessarily comes from a greater diversity of perspectives. The consequence is that development is more rapid; more rapid detection of problems and design bottlenecks, as more eyes pass over each of them, leads to more rapid repair and at a deeper level, as OS projects are better able to spot cross-platform synergies and eliminate redundant code than hierarchical, silo-ed development teams. This delivers a product that is more robust, leaner and thus more efficient.

There are also arguments in favour of open source as better able to serve the common good. The boundary between developer and consumer is blurred - the pro-sumer - so the product better reflects what is needed, rather than what the producer wishes to sell. Instead of relying on a small paid team that has a vested interest in following the rules, more independent participants have a primary interest in achieving the right answer. Similarly, there is less lag in the product development cycle as the 'focus group' function is inextricably intertwined with development rather than preceding it. The direct rewards of success are also more widely and equitably distributed - OS tends towards making a large number quite rich rather than a small number very rich.

[8] A classic example of this is the milk float, designed electric just after the war. Plenty of them are still around but the company isn't. They were so good that all dairies bought them; by the late fifties, all dairies had them and, also because they were so good, they didn't want any more, so the company went under.

However, the single most compelling reason for an OS strategy in this instance is not related to software and the usual OS arguments; it very specifically emerges from the situation that transport technology is in at the moment and the technological opportunities that have arisen. As discussed earlier, there is a window of opportunity to readdress vehicle architecture, which has not changed for over 100 years, during which the latter technical approach cannot yield a commercially viable vehicle but the former can.

It is crucial to make the most of this window of opportunity. This focuses the question on how to maximise the chance of successfully commercialising the Network Hybrid; we need to achieve significant penetration, as soon as possible, so we need to engage as much support as possible. In such circumstances, there is a positive interest in encouraging competition, as the more products and manufacturers that adopt these technical standards, the greater the chance of them becoming ubiquitous. Giving away the tools for others to create entrepreneurial opportunities is the best way in which to engage a massive and distributed support base. The free dissemination of IP also allows for much more rapid innovation and progress than a system whereby monopolies are granted for 20 years whenever any significant advance is made.

The scenario on which we are working allows for a Not-for-profit design foundation that grants licenses symmetrically to any manufacturer who is willing and able to comply with the license terms. There will be a nominal fee attached, pro rata to units produced, in order to maintain the foundation. This will employ a small but effective team, largely based on motorsport experience, of salaried engineers who perform a support function to the OS community, responsible for editing, coordinating and archiving the design work, as well as version control. As the community develops, the batons of decision-making authority are gradually passed out into the network.

In the free software movement, free refers to 'free speech, not free beer.' No community is going to devote themselves to lining someone else's pocket but there is no objection to making profit by adding value to the core technology, such as bundling applications with Linux and providing a user interface for non-techies. The value could equally be added by converting a CAD file into something in which a consumer could go shopping, so manufacturers would be conventional commercial entities. The potential of this model as an intellectual powerhouse is extraordinary. As New Scientist said, Linux can harness "a pool of creativity that Microsoft, for all its huge resources, will never be able to match". For Microsoft, read Ford or Toyota.

Alternative solutions and commonly cited objections

There are a few related areas not addressed in this paper, as they are not central to the exposition of the proposed solution. Nonetheless, questions are begged by the foregoing proposals and so I have added some short notes on the main issues about which perfectly valid questions can, and are, raised. These relate to the wisdom of so many simultaneous changes and the production of hydrogen. There are also alternative vehicle technologies, not without merits; although space precludes their discussion, the 21st century will surely have places for a greater diversity of niche solutions than the 20th. The distribution of hydrogen is also considered a barrier but, as shown by E4Tech and Imperial College in a major study [9] for Linde Gas last year, the costs of a hydrogen infrastructure are not huge, especially considering the importance of transport refuelling and the expected longevity of the hydrogen solution once adopted. The study forecasts a figure of €3.5bn by 2020 for a Europe-wide network.

One step at a time?

> "You think that because you understand one you understand two,
> because one and one makes two, but you must also understand and."
>
> Sufi proverb

[9] http://www.linde.de/international/web/linde/likelindeeng.nsf/DocByAlias/news_F82F09F7

A common question posed is whether it really is necessary to do all this at once [10]. Unequivocally, the answer is yes. The proposals take the methodology of Whole System Design, as applied to the Hypercar, into the realm of business strategy and their strength emerges from the synergy between them; few of these suggestions make any performance or economic sense on their own and are, to varying degrees, completely incompatible with the existing models. The combination allows the prospect of sustainable personal transport, reduced throughput of resources, cheaper transport for the consumer and yet a more profitable industry, giving higher quality employment and delivering more genuinely customised solutions to the consumer at a cheaper lifecycle cost.

The quickest and most practical way to achieve what is required is to do it all at once. Far from increasing risk, the simultaneous adoption of the many complementary strategies actually lowers barriers and reduces business risk; incremental adoption of these strategies simply would not work.

The concept of appropriate scale, which permeates many aspects of sustainable design, is equally relevant to business strategy and explains how Daimler-Chrysler make a greater loss per unit on the Smart than any other car in production, on sales of 120,000. What they do, they do spectacularly well, but this niche is simply too small for their model to service profitably. If a major manufacturer is to survive the disruption of shifting to a completely new vehicle technology, penetration must ramp up very quickly whereas OSCar expects the disruption to be a gradual process, starting from low-powered niches. This is a very different trajectory than that envisaged by the industry, which expects fuel cell cars to initially gain a market foothold in the premium sector.

Hydrogen generation

Without doubt, the aspect of a hydrogen economy with which thoughtful sceptics find they have the most leverage is the production of hydrogen. This is not to say that it is a weakness in the case for hydrogen but that it is the most difficult to explain, as it unifies energy and transport strategy - one of the most complex scenario planning exercises facing mankind - and the power of the hydrogen proposal emanates from the synergies between them. Without hydrogen as an integral element in our energy strategy, it is hard to justify it as a ubiquitous transport fuel; argued from a whole system perspective, it is a different story, but it is culturally counter-intuitive for us to let go of direct one-on-one comparisons. The flaw in such comparisons is that hydrogen is a many sided phenomenon and, in Venn diagram terms, it overlaps with electricity, petrol, batteries, hydro-electric storage, biogas and others, but is not directly equivalent to any of them.

The central questions are the energy required to generate the hydrogen and the potential displacement of emissions, such as when cracking natural gas centrally. Although hydrogen does not grow on trees, neither does electricity. Whilst sequestration is certainly not sustainable, CO2 is only displaced while we continue to use unsustainable sources of hydrogen; many sources are sustainable (these do not include cracking natural gas or nuclear-powered electrolysis); and as an interim measure, hydrogen from natural gas, at 70%, is much more efficient than electricity from natural gas, which is 39% efficient on average and 49% when best in class.

As in all things sustainable, we must start with efficiency measures to reduce hydrogen demand before we plan where to get it. We can certainly have a better quality of life with a fraction of today's energy consumption. At least a factor 5 improvement is possible with transport and many sectors can do better than that. As a source, electricity is a high value form of energy and electricity should be prioritised for those applications for which it is irreplaceable, only to be used to balance hydrogen supply. There are many low grade energy sources, such as gasification of bio-waste streams, from which hydrogen can be extracted economically but electricity cannot without further significant losses, and many others, such as algae that give off hydrogen, under serious development. Nobody is suggesting that hydrogen will replace electricity as our primary energy grid but by adopting hydrogen in a complementary role, we greatly expand the universe of potential fungible renewable energy sources.

[10] As mentioned in the introduction, this paper does not actually cover all aspects of OSCar's integrated strategy; this includes manufacturing strategy, distribution and governance.

If we then proceed to evaluate at the whole system level, simplistic costing as transport fuel does not give a true picture of the value of hydrogen to the whole system. Intermittent renewables can use hydrogen as a load-levelling medium, both to balance grid load and contribute to transport fuel. Furthermore, the waste in electrolysis is heat and if this is fed into a heat grid, should we count it as waste? Most industrial energy losses are in the form of low temperature heat and this would take a significant chunk out of buildings' energy demand.

As we are now accustomed to hearing, hydrogen is an energy carrier not a source. It also can be generated from any energy source, so it decouples transport fuel from energy sources. This means that, once we implement hydrogen as the primary road transport energy vector and invest in the infrastructure, we never need make a similar investment again; we can migrate from 100% 'brown' hydrogen to 100% 'green' hydrogen, incrementally and at any rate we choose, without further investment in infrastructure or the consumer being aware of the change. However, this also applies to vehicle technology; it becomes independent of energy source and we never need another disruption in vehicle architecture. This gives real flexibility in developing our energy strategies in a transition to a sustainable world and, with a combination of whole system design and energy efficiency, there is every reason to believe that we will be able to harvest ample energy. This flexibility maximises the range of options, both regionally and temporally, open to us in the light of technical, economic and political changes; and we need flexibility.