

Making Humans a Multi-Planetary Species

Elon Musk

Chief Executive Officer
SpaceX
Hawthorne, California.

By talking about the SpaceX Mars architecture, I want to make Mars seem possible—make it seem as though it is something that we can do in our lifetime. There really is a way that anyone could go if they wanted to.

WHY GO ANYWHERE?

I think there are really two fundamental paths. History is going to bifurcate along two directions. One path is we stay on Earth forever, and then there will be some eventual extinction event. I do not have an immediate doomsday prophecy, but eventually, history suggests, there will be some doomsday event.

The alternative is to become a space-bearing civilization and a multi-planetary species, which I hope you would agree is the right way to go.

So how do we figure out how to take you to Mars and create a self-sustaining city—a city that is not merely an outpost but which can become a planet in its own right, allowing us to become a truly multi-planetary species (see *Fig. 1*)?

WHY MARS?

Sometimes people wonder, “Well, what about other places in the solar system? Why Mars?” Our options for becoming a multi-planetary species within our solar system are limited. We have, in terms of nearby options, Venus, but Venus is a high-pressure—super-high-pressure—hot acid bath, so that would be a tricky one. Venus is not at all like the goddess. So, it would be really difficult to make things work on Venus.

Then, there is Mercury, but that is way too close to the sun. We could potentially go onto one of the moons of Jupiter or Saturn, but those are quite far out, much further from the sun, and much harder to get to.

It really only leaves us with one option if we want to become a multi-planetary civilization, and that is Mars. We could conceivably go to our moon, and I actually have nothing

against going to the moon, but I think it is challenging to become multi-planetary on the moon because it is much smaller than a planet. It does not have any atmosphere. It is not as resource-rich as Mars. It has got a 28-day day, whereas the Mars day is 24.5 hours. In general, Mars is far better-suited ultimately to scale up to be a self-sustaining civilization.

To give some comparison between the two planets, they are remarkably close in many ways (*Table 1*). In fact, we now believe that early Mars was a lot like Earth. In effect, if we could warm Mars up, we would once again have a thick atmosphere and liquid oceans.

Mars is about half as far again from the sun as Earth is, so it still has decent sunlight. It is a little cold, but we can warm it up. It has a very helpful atmosphere, which, being primarily CO₂ with some nitrogen and argon and a few other trace elements, means that we can grow plants on Mars just by compressing the atmosphere.

It would be quite fun to be on Mars because you would have gravity that is about 37% of that of Earth, so you would be able to lift heavy things and bound around. Furthermore, the day is remarkably close to that of Earth. We just need to change the populations because currently we have seven billion people on Earth and none on Mars.

FROM EARLY EXPLORATION TO A SELF-SUSTAINING CITY ON MARS

There has been a lot of great work by NASA and other organizations in the early exploration of Mars and understanding what Mars is like. Where could we land? What is the composition of the atmosphere? Where is there water or ice? We need to go from these early exploration missions to actually building a city.

The issue that we have today is that if you look at a Venn diagram, there is no intersection of sets—of people who want to go and those who can afford to go (*Fig. 2*). In fact, right now, you cannot go to Mars for infinite money.

Using traditional methods, taking an Apollo-style approach, an optimistic cost would be about \$10 billion per person. Taking the Apollo program as an example, the cost estimates are somewhere between \$100 and \$200 billion in

This paper is a summary of Elon Musk's presentation at the 67th International Astronautical Congress in Guadalajara, Mexico, September 26–30, 2016. In February 2017, SpaceX announced it will launch a crewed mission beyond the moon for two private customers in late 2018.

Used with permission from SpaceX.




Fig. 1. Artist's rendition of a Martian colonist.

current-year dollars, and we sent 12 people to the surface of the moon, which was an incredible thing—probably one of the greatest achievements of humanity.

However, that is a steep price to pay for a ticket. That is why these circles barely touch (Fig. 3). You cannot create a self-sustaining civilization if the ticket price is \$10 billion per person. What we need to do is to move those circles together (Fig. 4). If we can get the cost of moving to Mars to be roughly equivalent to a median house price in the United States, which is around \$200,000, then I think the probability of establishing

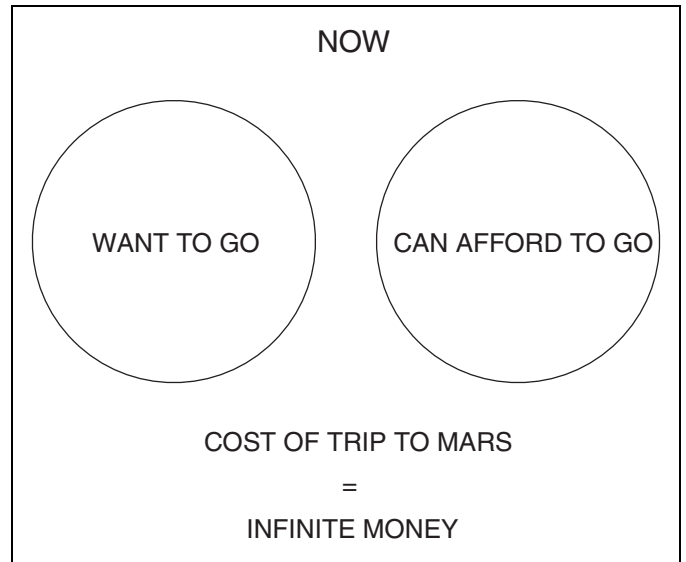


Fig. 2. Venn diagram of people who want to go to Mars versus those who can afford to go now.

a self-sustaining civilization is very high. I think it would almost certainly occur.

Not everyone would want to go. In fact, probably a relatively small number of people from Earth would want to go, but enough would want to go who could afford it for it to happen. People could also get sponsorship. It gets to the point where almost anyone, if they saved up and this was their goal,

Table 1. Characteristics of Earth and Mars

	EARTH	MARS
DIAMETER	12,756 km / 7,926 mi	6,792 km / 4,220 mi
AVERAGE DISTANCE FROM SUN	150,000,000 km / 93,000,000 mi	229,000,000 km / 142,000,000 mi
TEMPERATURE RANGE	-88C TO 58C / -126F TO 138F	-140C TO 30C / -285F TO 88F
ATMOSPHERIC COMPOSITION	78% N ₂ , 21% O ₂ , 1% OTHER	96% CO ₂ , < 2% Ar, < 2% N ₂ , < 1% OTHER
FORCE OF GRAVITY (WEIGHT)	100 lbs ON EARTH	38 lbs ON MARS (62.5% LESS GRAVITY)
DAY LENGTH	24 hrs	24 hrs 40 min
LAND MASS	148.9 MILLION km ²	144.8 MILLION km ² (97% OF EARTH)
PEOPLE	7 BILLION	0

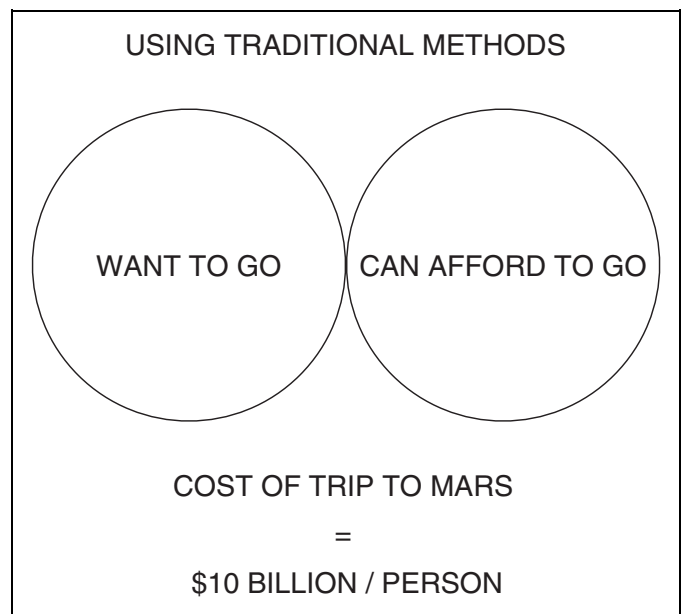


Fig. 3. Venn diagram of people who want to go to Mars versus those who can afford to go using traditional methods.

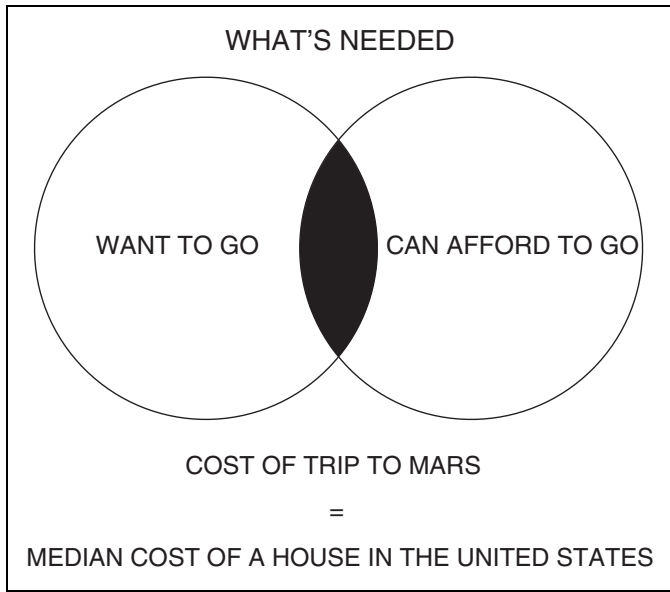


Fig. 4. Estimated price per ticket where there is overlap in a Venn diagram of people who want to go to Mars versus those who can afford to go.

could buy a ticket and move to Mars—and given that Mars would have a labor shortage for a long time, jobs would not be in short supply.

IMPROVING COST PER TON TO MARS BY FIVE MILLION PERCENT

It is a bit tricky because we have to figure out how to improve the cost of trips to Mars by five million percent. This translates to an improvement of approximately four-and-a-half orders of magnitude. This is not easy. It sounds virtually impossible, but there are ways to do it.

These are the key elements that are needed in order to achieve the four-and-a-half orders of magnitude improvement. Most of the improvement would come from full reusability—somewhere between two and two-and-a-half orders of magnitude. The other two orders of magnitude would come from refilling in orbit, propellant production on Mars, and choosing the right propellant.

Full reusability

To make Mars trips possible on a large enough scale to create a self-sustaining city, full reusability is essential. Full reusability is really the super-hard one. It is very difficult to achieve reusability even for an orbital system, and that challenge becomes substantially greater for a system that has to go to another planet.

You could use any form of transport as an example of the difference between reusability and expendability in aircraft. A

Table 2. Benefits of Refilling in Orbit

Not refilling in orbit would require a 3-stage vehicle at 5–10x the size and cost
Spreading the required lift capacity across multiple launches substantially reduces development costs and compresses schedule
Combined with reusability, refilling makes performance shortfalls an incremental rather than exponential cost increase

car, bicycle, horse, if they were single-use—almost no one would use them; it would be too expensive. However, with frequent flights, you can take an aircraft that costs \$90 million and buy a ticket on Southwest right now from Los Angeles to Vegas for \$43, including taxes. If it were single use, it would cost \$500,000 per flight. Right there, you can see an improvement of four orders of magnitude.

Now, this is harder—reusability does not apply quite as much to Mars because the number of times that you can reuse the spaceship pod of the system is less often because the Earth–Mars rendezvous only occurs every 26 months. Therefore, you get to use the spaceship part approximately every 2 years.

Refilling in orbit

You would get to use the booster and the tanker frequently. Therefore, it makes sense to load the spaceship into orbit with essentially tanks dry. If it has really big tanks that you use the booster and tanker to refill once in orbit, you can maximize the payload of the spaceship, so when it goes to Mars, you have a very large payload capability.

Hence, refilling in orbit is one of the essential elements of this (Table 2). Without refilling in orbit, you would have roughly a half order of magnitude impact on the cost. What that means is that each order of magnitude is a factor of 10. Therefore, not refilling in orbit would mean roughly a 500% increase in the cost per ticket.

It also allows us to build a smaller vehicle and lower the development cost, although this is still quite big. However, it would be much harder to build something that is 5–10 times the size.

Table 3. Benefits of Propellant Production on Mars

Allows reusability of the ship and enables people to return to Earth easily
Leverages resources readily available on Mars
Bringing return propellant requires approximately 5 times as much mass departing Earth

Furthermore, it reduces the sensitivity of the performance characteristics of the booster rocket and tanker. So, if there is a shortfall in the performance of any of the elements, you can make up for it by having one or two extra refilling trips to the spaceship. This is very important for reducing the susceptibility of the system to a performance shortfall.

Propellant production on Mars

Producing propellant on Mars is obviously also very important (Table 3). Again, if we did not do this, it would have at least a half order of magnitude increase in the cost of a trip. It would be pretty absurd to try to build a city on Mars if your spaceships just stayed on Mars and did not go back to Earth. You would have a massive graveyard of ships; you have to do something with them.

It would not really make sense to leave your spaceships on Mars; you would want to build a propellant plant on Mars and send the ships back. Mars happens to work out well for that because it has a CO₂ atmosphere, it has water-ice in the soil, and with H₂O and CO₂, you can produce methane (CH₄) and oxygen (O₂).

Right propellant

Picking the right propellant is also important. There are three main choices, and they each have their merits (Table 4). First, there is kerosene, or rocket propellant-grade kerosene, essentially a highly refined form of jet fuel. It helps keep the vehicle size small, but because it is a very specialized form of jet fuel, it is quite expensive. Its reusability potential is lower. It would be very difficult to make this on Mars

Table 4. Comparison of Kerosene, Hydrogen/Oxygen, and Deep-Cryo Methalox Propellants

	C ₁₂ H _{22.4} / O ₂ KEROSENE	H ₂ / O ₂ HYDROGEN/ OXYGEN	CH ₄ / O ₂ DEEP-CRYO METHALOX
VEHICLE SIZE	●	●	●
COST OF PROP	●	●	●
REUSABILITY	●	●	●
MARS PROPELLANT PRODUCTION	×	●	●
PROPELLANT TRANSFER	●	●	●

● GOOD
 ● OK
 ● BAD
 × VERY BAD

because there is no oil. Propellant transfer is pretty good but not great.

Hydrogen, although it has a high specific impulse, is very expensive, and it is incredibly difficult to keep from boiling off because liquid hydrogen is very close to absolute zero as a liquid. Therefore, the installation required is tremendous, and the energy cost on Mars of producing and storing hydrogen would be very high.

Therefore, when we looked at the overall system optimization, it was clear that methane was the clear winner. Methane would require from 50% to 60% of the energy on Mars to refill propellant using the propellant depot, and the technical challenges are a lot easier. We therefore think methane is better almost across the board.

We started off initially thinking that hydrogen would make sense, but ultimately came to the conclusion that the best way to optimize the cost-per-unit mass to Mars and back is to use an all-methane system—or technically, deep-cryo methalox.

Whatever system is designed, whether by SpaceX or someone else, these are the four features that need to be addressed in order for the system really to achieve a low cost per ton to the surface of Mars.

SYSTEM ARCHITECTURE

Figure 5 describes the overall system (for a full simulation, see www.spacex.com/mars). The rocket booster and the spaceship take off and launch the spaceship into orbit. The rocket booster then comes back quite quickly, within about 20 minutes. So, it can actually launch the tanker version of the spacecraft, which is essentially the same as the spaceship but filling up the unpressurized and pressurized cargo areas with propellant tanks. This also helps lower the development cost, which obviously will not be small.

Then, the propellant tanker goes up anywhere from three to five times to fill the tanks of the spaceship in orbit. Once the tanks are full, the cargo has been transferred, and we reach the Mars rendezvous timing, which is roughly every 26 months, that is when the ship would depart.

Over time, there would be many spaceships. You would ultimately have upwards of 1,000 or more spaceships waiting in orbit. Hence, the Mars Colonial fleet would depart en masse.

It makes sense to load the spaceships into orbit because you have got 2 years to do so, and then you can make frequent use of the booster and the tanker to get really heavy reuse out of those. With the spaceship, you get less reuse because you have to consider how long it is going to last—maybe 30 years, which might be perhaps 12–15 flights of the spaceship at most.




Fig. 5. System architecture. Targeted reuse per vehicle: 1,000 uses per booster, 100 per tanker, 12 per ship.

Therefore, you really want to maximize the cargo of the spaceship and reuse the booster and the tanker as much as possible. Hence, the ship goes to Mars, gets replenished, and then returns to Earth.

This ship will be relatively small compared with the Mars interplanetary ships of the future. However, it needs to fit 100 people or thereabouts in the pressurized section, carry the luggage and all of the unpressurized cargo to build propellant plants, and to build everything from iron foundries to pizza joints to you name it—we need to carry a lot of cargo.

The threshold for a self-sustaining city on Mars or a civilization would be a million people. If you can only go every 2 years and if you have 100 people per ship, that is 10,000 trips. Therefore, at least 100 people per trip is the right order of magnitude, and we may end up expanding the crew section and ultimately taking more like 200 or more people per flight in order to reduce the cost per person.

However, 10,000 flights is a lot of flights, so ultimately you would really want in the order of 1,000 ships. It would take a while to build up to 1,000 ships. How long it would take to reach that million-person threshold, from the point at which the first ship goes to Mars would probably be somewhere between 20 and 50 total Mars rendezvous—so it would take 40–100 years to achieve a fully self-sustaining civilization on Mars.

VEHICLE DESIGN AND PERFORMANCE

Figure 6 is a cross-section of the ship. In some ways, it is not that complicated. It is made primarily of an advanced carbon fiber. The carbon-fiber part is tricky when dealing with deep cryogenics and trying to achieve both liquid and gas impermeability and not have gaps occur due to cracking or pressurization that would make the carbon fiber leaky. Hence, it is a fairly significant technical challenge to make deeply cryogenic tanks out of carbon fiber. It is only recently that carbon-fiber technology has reached the point where we can do this without having to create a liner on the inside of the tanks, which would add mass and complexity.

This is particularly tricky for the pressurization of the hot gases. This is likely to be autogenously pressurized, which means that we gasify the fuel and the oxygen through heat exchanges in the engine and use that to pressurize the tanks. So, we gasify the methane and use that to pressurize the fuel tank, and we gasify the oxygen and use that to pressurize the oxygen tank.




Fig. 6. Cross-section of the ship. Carbon fiber primary structure; densified CH₄/O₂ propellant; autogenous pressurization.

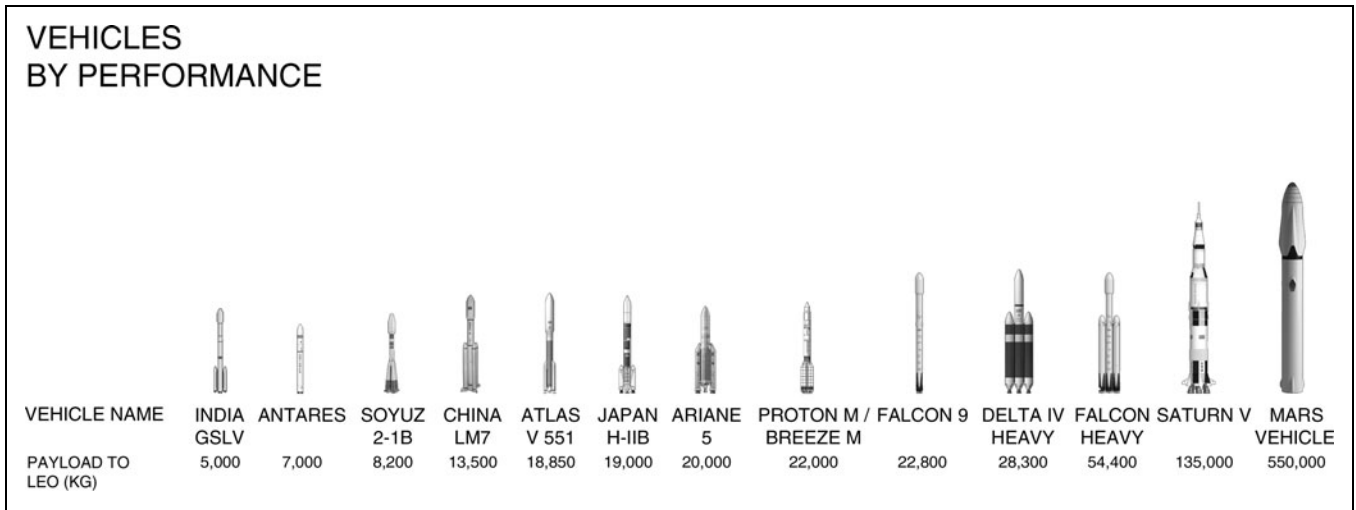


Fig. 7. Current and historic vehicles by performance.

This is a much simpler system than what we have with Falcon where we use helium for pressurization and nitrogen for gas thrusters. In this case, we would autogenously pressurize and then use gaseous methane and oxygen for the control thrusters. Hence, you really only need two ingredients for this, as opposed to four in the case of Falcon 9, or five if you consider the ignition liquid. In this case, we would use spark ignition.

Figure 7 gives you a sense of vehicles by performance, current and historic. For expendable mode, the vehicle that we were proposing would do about 550 tons and about 300 tons in reusable mode. That compares to the Saturn V max capability of 135 tons.

Figure 8 gives a better sense of things. The dark gray bars show the performance of the vehicle, the payload to orbit of

the vehicle. What it represents is the size efficiency of the vehicle. With most rockets, including ours, that are currently flying, the performance bar is only a small percentage of the actual size of the rocket.

However, with the interplanetary system, which will initially be used for Mars, we believe we have improved the design performance massively. It is the first time a rocket performance bar will actually exceed the physical size of the rocket.

Figure 9 gives you a more direct comparison. The thrust level is enormous. We are talking about a lift-off thrust of 13,000 tons, so it will be quite tectonic when it takes off. However, it does fit on Pad 39A, which NASA has been kind enough to allow us to use because they oversized the pad in doing Saturn V. As a result, we can use a much larger vehicle on that same launchpad.




Fig. 8. Payload to orbit of current and historic vehicles.




Fig. 9. Comparison of Mars vehicle and Saturn V.

In the future, we expect to add additional launch locations, probably adding one on the south coast of Texas, but this gives you a sense of the relative capability.

However, these vehicles have very different purposes. This is really intended to carry huge numbers of people, ultimately millions of tons of cargo to Mars. Therefore, you really need something quite large in order to do that.

RAPTOR ENGINE

We started the development with what are probably the two most difficult key elements of the design of the inter-planetary spaceship, the engine and rocket booster. The Raptor engine is going to be the highest chamber pressure engine of any kind ever built, and probably the highest thrust-to-weight (Fig. 10).





Fig. 10. Characteristics of the Raptor engine.

It is a full-flow staged combustion engine, which maximizes the theoretical momentum that you can get out of a given source fuel and oxidizer. We subcool the oxygen and methane to densify it. Compared with when used close to their *boiling* points in most rockets, in our case, we load the propellants close to their *freezing* point. That can result in a density improvement of around 10%–12%, which makes an enormous difference in the actual result of the rocket. It gets rid of any cavitation risk for the turbo pumps, and it makes it easier to feed a high-pressure turbo pump if you have very cold propellant.

One of the keys here, though, is the vacuum version of the Raptor having a 382-second ISP. This is critical to the whole Mars mission and we are confident we can get to that number or at least within a few seconds of that number, ultimately maybe even exceeding it slightly.

ROCKET BOOSTER

In many ways, the rocket booster is really a scaled-up version of the Falcon 9 booster (*Fig. 11*). There are a lot of similarities, such as the grid fins and clustering a lot of engines at the base. The big differences are that the primary structure is an advanced form of carbon fiber as opposed to aluminum lithium, we use autogenous pressurization, and we get rid of the helium and the nitrogen.



Length	77.5 m
Diameter	12 m
Dry Mass	275 t
Propellant Mass	6,700 t
Raptor Engines	42
Sea Level Thrust	128 MN
Vacuum Thrust	138 MN

Booster accelerates ship to staging velocity, traveling 8,650 km/h (5,375 mph) at separation

Booster returns to landing site, using 7% of total booster prop load for boostback burn and landing

Grid fins guide rocket back through atmosphere to precision landing

Fig. 11. Characteristics of the rocket booster.




Fig. 12. Configuration of the Raptor engines within the rocket booster.

Each rocket booster uses 42 Raptor engines (*Fig. 12*). It is a lot of engines, but with Falcon Heavy, which should launch early next year, there are 27 engines on the base. Therefore, we have considerable experience with a large number of engines. It also gives us redundancy so that if some of the engines fail, you can still continue the mission and everything will be fine.

However, the main job of the booster is to accelerate the spaceship to around 8,500 km/h. For those who are less familiar with orbital dynamics, it is all about velocity and not about height.

In the case of other planets, though, which have a gravity well that is not as deep, such as Mars, the moons of Jupiter, conceivably one day maybe even Venus—well, Venus will be a little trickier—but for most of the solar system, you only need the spaceship. You do not need the booster if you have a lower gravity well. Therefore, no booster is needed on the moon or Mars or any of the moons of Jupiter or Pluto. The booster is just there for heavy gravity wells.

We have also been able to optimize the propellant needed for boost back and landing to get it down to about 7% of the lift-off propellant load. With some optimization, maybe we can get it down to about 6%.

We are also now getting quite comfortable with the accuracy of the landing. If you have been watching the Falcon 9 landings, you will see that they are getting increasingly closer to the bull’s eye. In particular, with the addition of maneuvering thrusters, we think we can actually put the booster right back on the launch stand. Then, those fins at the base are essentially centering features to take out any minor position mismatch at the launch site.

So, that is what it looks like at the base. We think we only need to gimbal or steer the center cluster of engines. There are seven engines in the center cluster. Those would be the

ones that move for steering the rocket, and the other ones would be fixed in position. We can max out the number of engines because we do not have to leave any room for gimbaling or moving the engines. This is all designed so that you could actually lose multiple engines, even at lift-off or anywhere in flight, and still continue the mission safely.

INTERPLANETARY SPACESHIP

For the spaceship itself, in the top, we have the pressurized compartment. Then, beneath that is where we would have the unpressurized cargo, which would be really flat-packed—a very dense format. Below that is the liquid oxygen tank (Fig. 13).

The liquid oxygen tank is probably the hardest piece of this whole vehicle because it must handle propellant at the coldest level and the tanks themselves actually form the airframe. The airframe structure and the tank structure are combined, as is the case in all modern rockets. In aircraft, for example, the wing is really a fuel tank in the shape of a wing.

The oxygen tank has to take the thrust loads of ascent and the loads of reentry, and then it has to be impermeable to gaseous oxygen, which is tricky, and nonreactive to gaseous oxygen. Therefore, that is the most difficult piece of the spaceship itself, which is also why we started on that element.

Below the oxygen tank is the fuel tank, and then the engines are mounted directly to the thrust cone on the base. There are six of the high-efficiency vacuum engines around the perimeter, and those do not gimbal. There are three of the sea-level versions of the engine, which do gimbal and provide the steering, although we can do some amount of steering if you are in space with differential thrust on the outside engines.

The net effect is a cargo to Mars of up to 450 tons, depending upon how many refills you do with the tanker. The goal is at least 100 passengers per ship, although ultimately, we will probably see that number grow to 200 or more.

Depending upon which Earth-Mars rendezvous you are aiming for, the trip time at 6 km/s departure velocity can be as low as 80 days (Fig. 14).

Over time, we would improve that, and, eventually, I suspect that you would see Mars transit times of as little as 30 days in the more distant future. It is fairly manageable, considering the trips that people used to do in the old days where sailing voyages would take 6 months or more.

On arrival, the heat-shield technology is extremely important (Fig. 15). We have been refining the heat-shield technology using our Dragon spacecraft, and we are on version 3 of PICA, which is a phenolic-impregnated carbon ablator, and it is getting more robust with each new version,

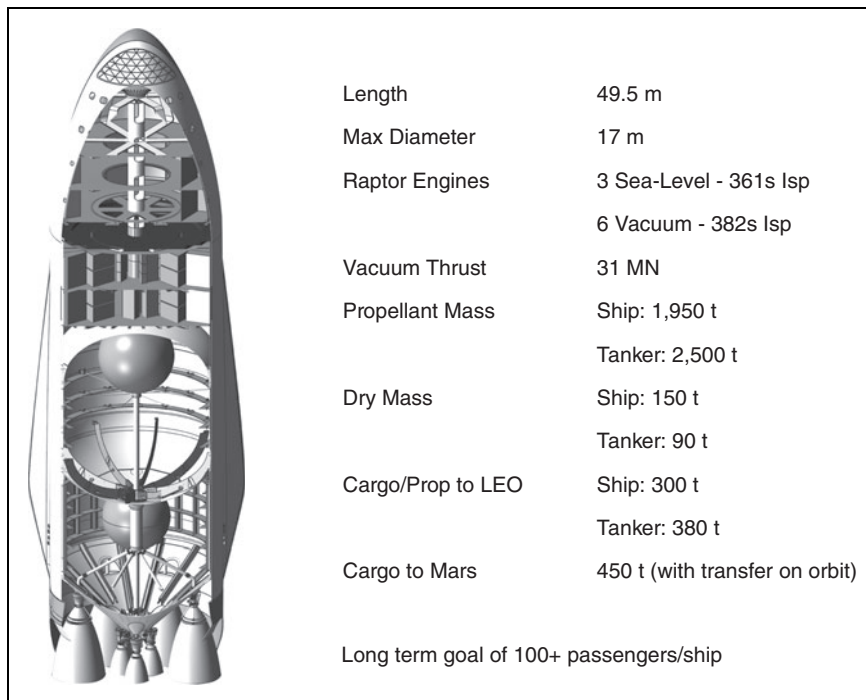


Fig. 13. Characteristics of the interplanetary spaceship.

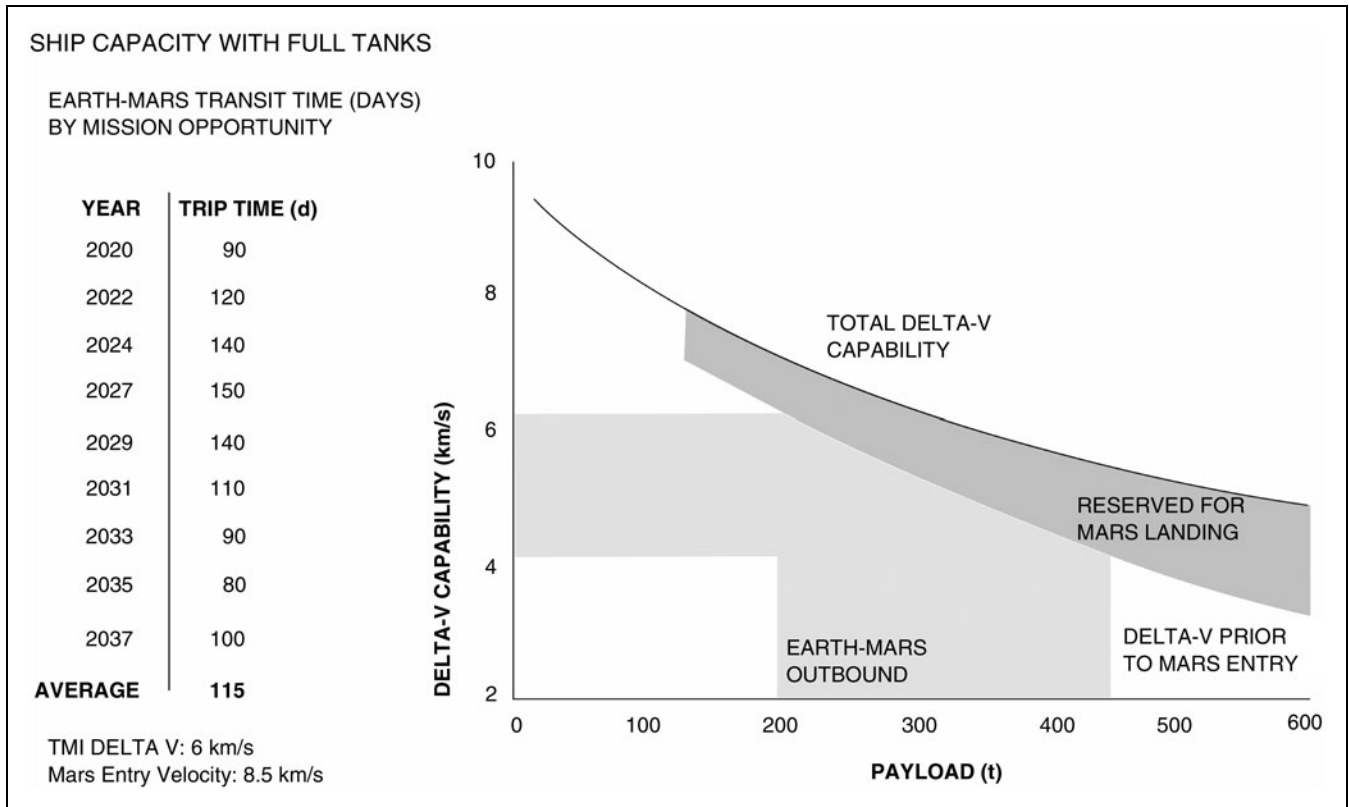


Fig. 14. Earth–Mars transit time (in days).

with less ablation, more resistance, and less need for refurbishment.

The heat shield is basically a giant brake pad. It is a matter of how good you can make that brake pad against extreme re-entry conditions, minimize the cost of refurbishment, and

make it so that you could have many flights with no refurbishment at all.

I want to give you a sense of what it would feel like to actually be in the spaceship. In order to make it appealing and increase that portion of the Venn diagram where people

ARRIVAL

From interplanetary space, the ship enters the atmosphere, either capturing into orbit or proceeding directly to landing

Aerodynamic forces provide the majority of the deceleration, then 3 center Raptor engines perform the final landing burn

Using its aerodynamic lift capability and advanced heat shield materials, the ship can decelerate from entry velocities in excess of 8.5 km/s at Mars and 12.5 km/s at Earth

G-forces (Earth-referenced) during entry are approximately 4-6 g's at Mars and 2-3 g's at Earth

Heating is within the capabilities of the PICA-family of heat shield materials used on our Dragon spacecraft

PICA 3.0 advancements for Dragon 2 enhance our ability to use the heat shield many times with minimal maintenance

Fig. 15. Heat-shield technology for arrival.

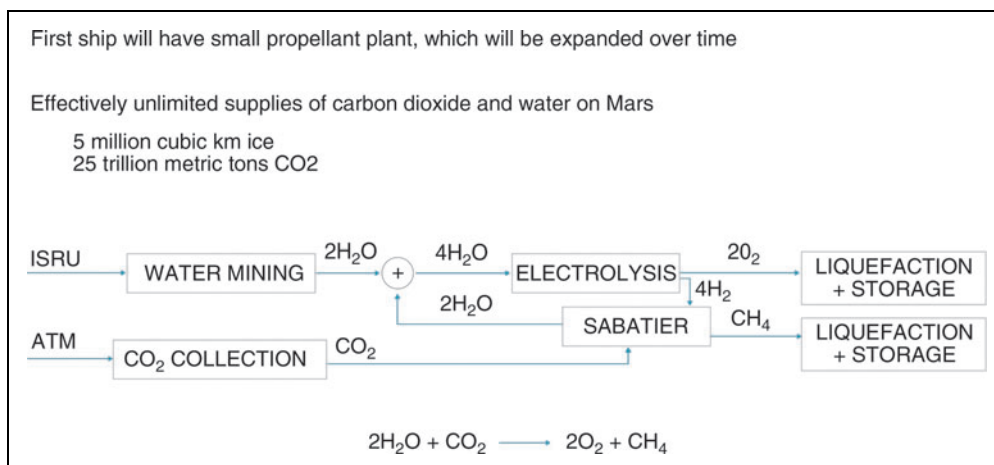


Fig. 16. Propellant production on Mars.

actually want to go, it has got to be really fun and exciting—it cannot feel cramped or boring. Therefore, the crew compartment or the occupant compartment is set up so that you can do zero-gravity games—you can float around. There will be movies, lecture halls, cabins, and a restaurant. It will be really fun to go. You are going to have a great time!

PROPELLANT PLANT

The ingredients are there on Mars to create a propellant plant with relative ease because the atmosphere is primarily CO₂, and water-ice is almost everywhere. There is CO₂ plus H₂O to make methane, CH₄, and oxygen, O₂, using the Sabatier reaction. The trickiest thing really is the energy source, which can be done with a large field of solar panels (Fig. 16).

COST PER TRIP

The key is making this affordable to almost anyone who wants to go (Fig. 17). Based on this architecture, assuming optimization over time, we are looking at a cost per ticket of <\$200,000, maybe as little as \$100,000 over time, depending upon how much mass a person takes.

Right now, we are estimating about \$140,000 per ton for the trips to Mars. If a person plus their luggage is less than that, taking into account food consumption and life support, the cost of moving to Mars could ultimately drop below \$100,000.

Obviously, it is going to be a challenge to fund this whole endeavor. We expect to generate a pretty decent net

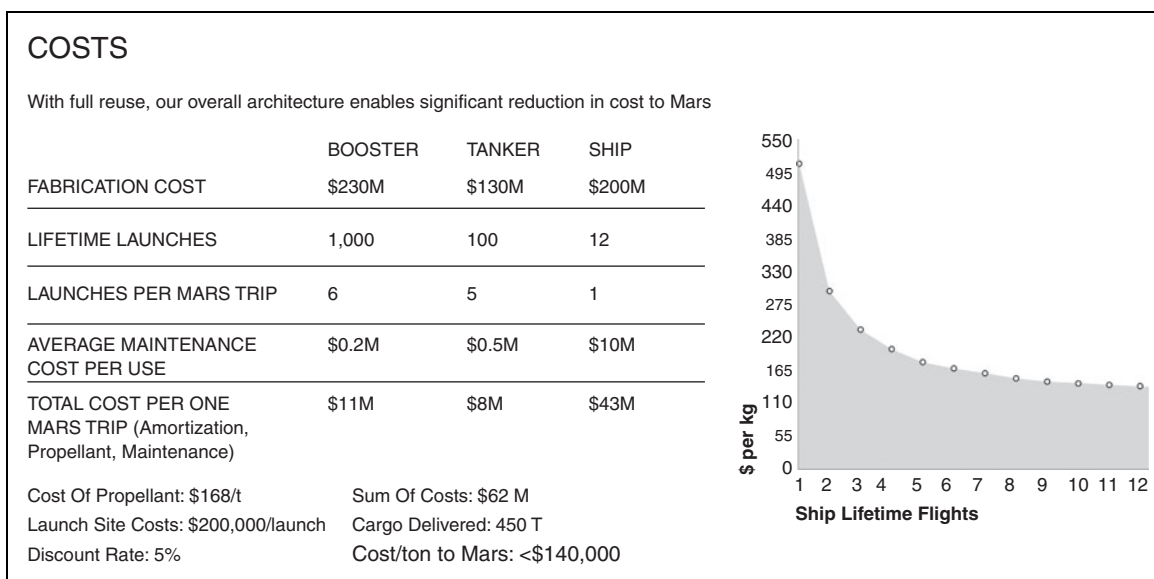


Fig. 17. Cost per ton to Mars.

cash flow from launching lots of satellites and servicing the space station for NASA, transferring cargo to and from the space station.

There are also many people in the private sector who are interested in helping to fund a base on Mars, and perhaps there will be interest on the government sector side to do that too. Ultimately, this is going to be a huge public-private partnership.

Right now, we are just trying to make as much progress as we can with the resources that we have available and to keep the ball moving forward. As we show that this is possible and that this dream is real—it is not just a dream, it is something that can be made real—the support will snowball over time.

I should also add that the main reason I am personally accumulating assets is in order to fund this. I really do not have any other motivation for personally accumulating assets except to be able to make the biggest contribution I can to making life multi-planetary.

TIMELINES

In 2002, SpaceX basically consisted of carpet and a mari-achi band. That was it. I thought we had maybe a 10% chance of doing anything—of even getting a rocket to orbit, let alone getting beyond that and taking Mars seriously.

However, I came to the conclusion that if there were no new entrants into the space arena with a strong ideological motivation, then it did not seem as if we were on a trajectory to ever be a space-based civilization and be out there among the stars.

In 1969, we were able to go to the moon, and the space shuttle could get to low Earth orbit. Then the space shuttle was retired. However, that trend line is down to zero. What many people do not appreciate is that technology does not automatically improve; it only improves if a lot of really strong engineering talent is applied to the problem. There are many examples in history where civilizations have reached a certain technology level, fallen well below that, and then recovered only millennia later.

We went from 2002, where we basically were clueless, and we built the smallest, useful orbital rocket that we could think of with Falcon 1, which would deliver half a ton to orbit. Four years later, we developed the first vehicle. We developed the main engine, the upper-stage engine, the airframes, the fairing, and the launch system, and we had our first attempt at launch in 2006, which failed. It lasted about 60 seconds, unfortunately.

However, 2006, 4 years after starting, was also when we got our first NASA contract (*Fig. 18*). I am incredibly grateful to




Fig. 18. SpaceX milestones.

MUSK

NASA for supporting SpaceX, despite the fact that our rocket crashed. I am NASA's biggest fan. Thank you very much to the people who had the faith to do that.

Finally, the fourth launch of Falcon 1 worked in 2008 (Fig. 18). We were really down to our last pennies. In fact, I only thought I had enough money for three launches, and the first three failed! We were able to scrape together enough to just make it and do a fourth launch, and, thank goodness—that fourth launch succeeded in 2008. That was a lot of pain.

The end of 2008 is also when NASA awarded us the first major operational contract, which was for resupplying cargo to the space station and bringing cargo back. A couple of years later, we did the first launch of Falcon 9, version 1, and that had about a 10-ton-to-orbit capability, which was about 20 times the capability of Falcon 1 (Fig. 18). It was also assigned to carry our Dragon spacecraft.

It was in 2012 when we delivered and returned cargo from the space station. In 2013, we started doing vertical take-off and landing tests for the first time (Fig. 18). Then, in 2014, we were able to have the first orbital booster do a soft landing in the ocean. The landing was soft, it fell over, and it exploded. However, for 7 seconds, it was good. We also improved the capability of the vehicle from 10 tons to about 13 tons to LEO. December 2015, was definitely one of the best moments of my life: the rocket booster came back and landed at Cape Canaveral (Fig. 18). That really showed that we could bring an orbital-class booster back from a very high velocity, all the way to the launch site, and land it safely with almost no refurbishment required for reflight. If things go well, we are hoping to reflly one of the landed boosters in a few months.

In 2016, we also demonstrated landing on a ship (Fig. 18), which is important for both the very high-velocity geosynchronous missions and for the reusability of Falcon 9 because about roughly a quarter of our missions service the space station. There are a few other lower-orbit missions, but probably 60% of our missions are commercial GEO missions. These high-velocity missions need to land on a ship out at sea. They do not have enough propellant onboard to boost back to the launch site.

FUTURE

Figure 19 shows the future—next steps. We were intentionally fuzzy about this timeline. However, we are going to try to make as much progress as we can on a very constrained budget, on the elements of the interplanetary transport booster and spaceship. Hopefully, we will be able to complete the first development spaceship in maybe about 4 years, and we will start doing suborbital flights with that.

It has enough capability that you could possibly go to orbit if you limit the amount of cargo on the spaceship. You would have to really strip it down, but in tanker form, it could definitely get to orbit. It cannot get back, but it can get to orbit.

Maybe there is some market for the really fast transport of things around the world, provided we can land somewhere where noise is not a super-big deal because rockets are very noisy. We could transport cargo to anywhere on Earth in 45 minutes at the most. Hence, most places on Earth would be 20–25 minutes away. If we had a floating platform off the coast of New York, 20–30 miles out, you could go from New York to Tokyo in 25 minutes and across the Atlantic in 10 minutes. Most of your time would be spent getting to the ship,

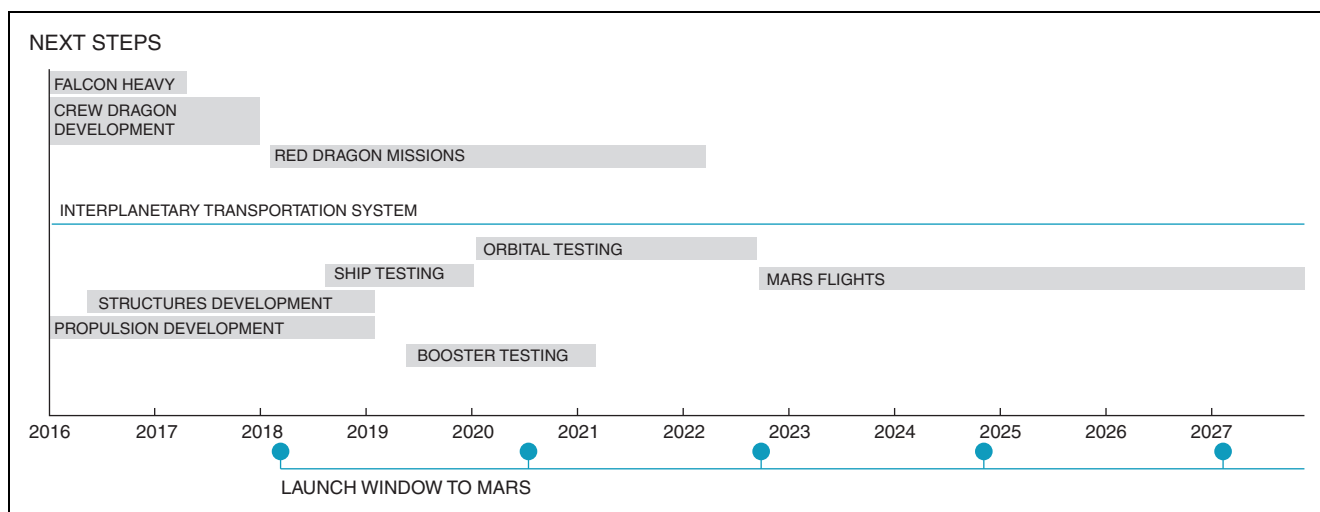


Fig. 19. Next steps in the development of the interplanetary transport system.

and then it would be very quick after that. Therefore, there are some intriguing possibilities there, although we are not counting on that.

Then, there is the development of the booster. The booster part is relatively straightforward because it amounts to a scaling up of the Falcon 9 booster. So, we do not see that there will be many showstoppers there.

Then it will be a case of trying to put it all together and make this actually work for Mars. If things go super-well, it might be in the 10-year timeframe, but I do not want to say that is when it will occur. There is a huge amount of risk. It is going to cost a lot. There is a good chance we will not succeed, but we are going to do our best and try to make as much progress as possible.

RED DRAGON

We are going to try to send something to Mars on every Mars rendezvous from this point on. We plan to send Dragon 2, which is a propulsive lander, to Mars in a couple of years, and then probably do another Dragon mission in 2020 (Fig. 20).

We want to establish a steady cadence, so that there is always a flight leaving, just like there is a train leaving the station. With every Mars rendezvous, we will be sending at least a Dragon to Mars and ultimately the big spaceship. If there are people who are interested in putting payloads on Dragon, you know you can count on a ship that is going to transport something on the order of at least two or three tons of useful payloads to the surface of Mars.

That is part of the reason why we designed Dragon 2 to be a propulsive lander. As a propulsive lander, you can go anywhere in the solar system. You could go to the moon. You could go to, well, anywhere! However, if something relies on wings, you can pretty much only land on Earth because you

need a runway, and most places do not have a runway. As for parachutes, you cannot use those anywhere that does not have a dense atmosphere. However, propulsive works anywhere. Therefore, we should be able to land Dragon on any solid or liquid surface in the solar system.

RAPTOR FIRING

I was really excited to see all our Raptor engines firing (Fig. 21). The Raptor is a really tricky engine. It is a lot trickier than a Merlin because it is a full-flow stage combustion, with much higher pressure. I am amazed that it did not blow up on the first firing, but fortunately it was good.

Part of the reason for making the engine small, although it has three times the thrust of a Merlin, it is actually only about the same size as a Merlin engine because it has three times the operating pressure. That means we can use many of the production techniques that we honed with Merlin.

We are currently producing Merlin engines at almost 300 per year. Therefore, we understand how to make rocket engines in volume. Hence, even though the Mars vehicle uses 42 on the base and nine on the upper stage—so we have 51 engines to make—that is well within our production capabilities for Merlin. This is a similarly sized engine to Merlin, except for the expansion ratio. We therefore feel really comfortable about being able to make this engine in volume at a price that does not break our budget.

CARBON-FIBER TANK

We also wanted to make progress on the primary structure. As I mentioned, this is really a very difficult thing to make out of carbon fiber, even though carbon fiber has incredible strength to weight. When you then want to put super-cold




Fig. 20. Red Dragon mission objectives.



Fig. 21. Raptor engine firing.

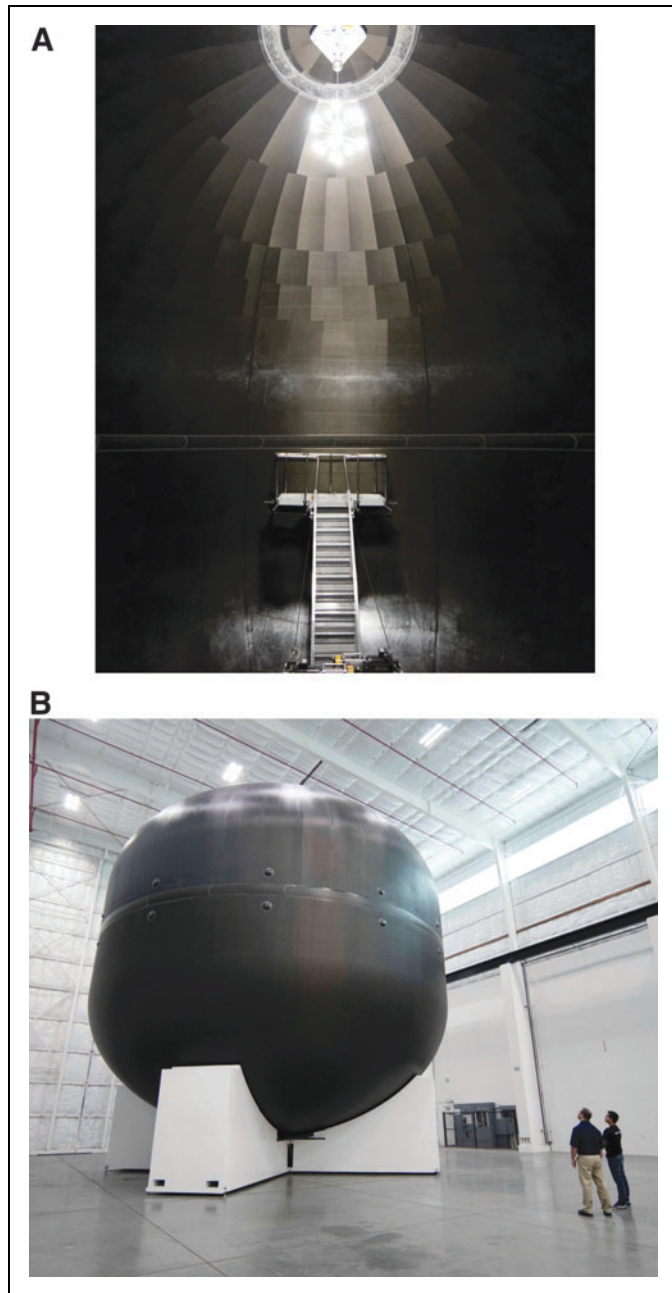


Fig. 22. (A) Interior and (B) exterior views of the carbon fiber tank.

liquid oxygen and liquid methane, particularly liquid oxygen, in the tank, it is subject to cracking and leaking.

The sheer scale of it is also challenging because you have to lay out the carbon fiber in exactly the right way on a huge mold, and you have to cure that mold at temperature. It is just really hard to make large carbon-fiber structures that could do all of those things and carry incredible loads.



Fig. 23. Flyby of Jupiter.

That was the other thing we wanted to focus on: the first development tank for the Mars spaceship. This is really the hardest part of the spaceship. The other pieces we have a pretty good handle on, but this was the trickiest one so we wanted to tackle it first.

This was a massive achievement. Huge congratulations are due to the team that worked on it. We managed to build the first tank, and the initial test with the cryogenic propellant actually looks quite positive. We have not seen any leaks or major issues.

Figure 22 is what the tank looks like on the inside. You get a real sense of just how big this tank is. It is completely smooth on the inside, but the way that the carbon fiber applies, lays up, and reflects the light makes it look multifaceted.

BEYOND MARS

What about beyond Mars? As we thought about the system and the reason we call it a system—because generally, I do not



Fig. 24. Propellant depot on Enceladus.



Fig. 25. Propellant depot on Europa.

like calling things “systems,” as everything is a system, including your dog. However, it is actually more than a vehicle. There is obviously the rocket booster, the spaceship, the tanker and the propellant plant, and the in situ propellant production.

If you have all four of these elements, you can go anywhere in the solar system by planet hopping or moon hopping. By establishing a propellant depot on the asteroid belt or on one of the moons of Jupiter, you can make flights from Mars to Jupiter. In fact, even without a propellant depot at Mars, you can do a flyby of Jupiter (*Fig. 23*).




Fig. 26. Flyby of Saturn.

However, by establishing a propellant depot, say on Enceladus (*Fig. 24*) or Europa (*Fig. 25*), and then establishing another one on Titan, Saturn’s moon, and then perhaps another one further out on Pluto or elsewhere in the solar system, this system really gives you the freedom to go anywhere you want in the greater solar system (*Fig. 26*).

Therefore, you could travel out to the Kuiper Belt, to the Oort cloud. I would not recommend this for interstellar journeys, but this basic system—provided we have filling stations along the way—means full access to the entire greater solar system.