COMPASSION8INNOVATION

FIRST EDITION · 2026

A STUDY BOOK IN SPACE SYSTEMS ENGINEERING

PART 2 THE TWELVE SYSTEMS

LUNAR HABITAT ACT

SITE **89°41′ S** | SHACKLETON RIM | CREW **4** | ROTATION **180 D** | LANDED **48.3 T**

Made from the ground it stands on

WRITTEN BY

Vaisakh Ramakrishna Sarma · Joydeep Hazra

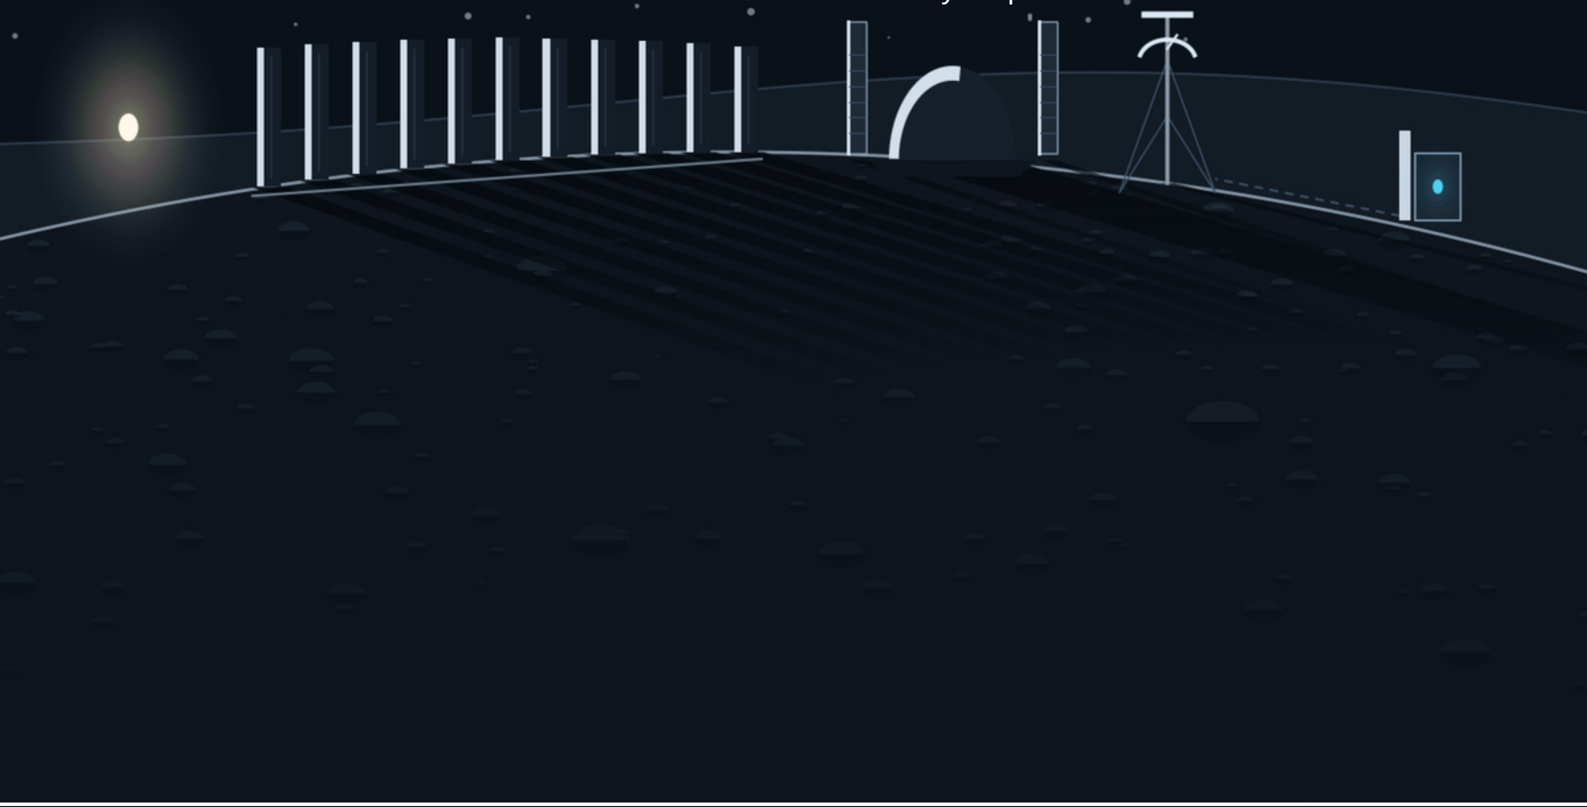

COMPASSION8INNOVATION
SCIENCE AND ENGINEERING TITLES

FIRST EDITION 2026

A STUDY BOOK IN SPACE SYSTEMS ENGINEERING

LUNAR HABITAT ACT

Five pressure vessels on skirts of compacted regolith, around a central core. The purpose of this outpost is not shelter. It is to keep four people alive long enough to build the Moon's first infrastructure.

Lunar Habitat ACT

A study book in space systems engineering

WRITTEN BY	Vaisakh Ramakrishna Sarma and Joydeep Hazra
PUBLISHED BY	Compassion8Innovation
EDITION	First edition, September 2026
SUBJECT	Space systems engineering, physics, design and technology
READERSHIP	Any reader
SOURCE DESIGN	Lunar Habitat ACT, the Autonomous Construction Testbed Concept definition, revision A

Every checkable figure in this book was put against the published record: NASA technical reports and standards, the Lunar Orbiter Laser Altimeter and Diviner data products, the Lunar Sourcebook, and peer reviewed papers. Where a source disagrees with the design, this book prints the source figure and says where the difference lies. Where nothing could be found to check a figure against, it says that too. The back matter carries every source and a full list of what the check found.

The cover shows the outpost looking north across the landing pad, from the design's own visualisation. Maps of the south polar region are drawn from published crater positions and are schematic: they are meant to show where things lie in relation to each other, not to be navigated by.

Contents

PART TWO

The twelve systems

A base is not a building. It is twelve machines that happen to share a hull, and each makes something another one needs. This part takes them one at a time: what the Moon does to you, what the physics allows, what this design chose, and where its numbers hold and where they do not. Each chapter ends with a mass and a power figure, as though the system could be bought alone. None can, and Part three puts them back together.

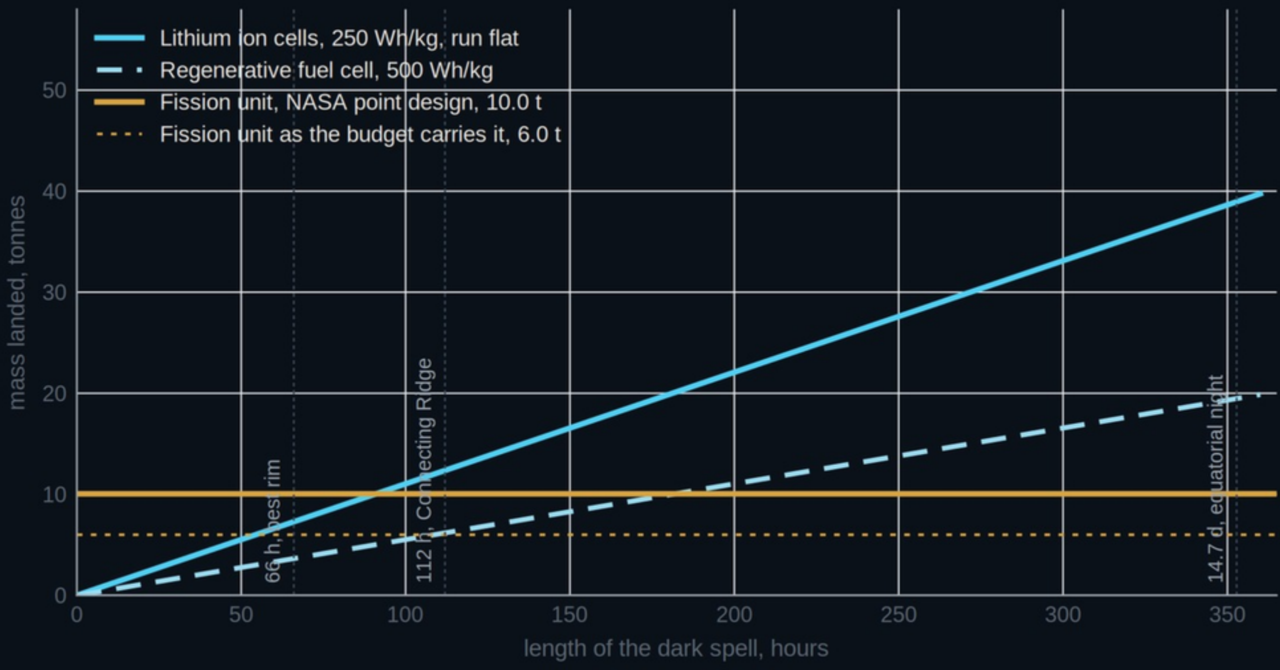

Power

The load on this base never falls below 27.6 kilowatts, and almost none of it can be switched off. Every power design for the Moon is an argument about what happens when the Sun goes.

WHAT THIS CHAPTER WILL TEACH YOU

1. Build a load table for a base, and tell a load that can be shed from one that cannot.
2. Explain why the night at this site is set by the terrain and not by the calendar.
3. Work out the battery mass a given dark spell demands, and see where the idea fails.
4. Say why this base carries a reactor when mass alone does not choose one.

48 kW	**210 kWh**	**40 kWe**	**1.5 to 5 days**
ARRAY PEAK	BATTERY INSTALLED	REACTOR CLASS	LONGEST DARK SPELL

5.1 Begin with the load, not with the Sun

The greenhouse draws 14 kilowatts. So does the oxygen plant. One of them is a chemical works that takes rock apart to get at the oxygen locked inside it, and the other is a room of 54 cubic metres with salad growing in racks. That they draw the same is the first surprise of this chapter, and it repays a moment's thought.

The plants are the only thing on this base for which sunlight has to be manufactured, and manufacturing sunlight is expensive. It is also the most movable load on the site, because a lettuce does not care whether its day lines up with the Sun's. Almost every concept study for a Moon base opens by announcing how much power it will generate. That is the wrong end of the problem. Generation is a choice; the load is not. Before anybody argues about panels or reactors, somebody has to write down what the base actually consumes and be honest about how much of it can be turned off when supply is short.

Figure 4.3 in the last chapter is that list, and it is worth reading again with a different question in mind. Read down the notes rather than the bars. Life support runs continuously and cannot be shed, because the crew keep breathing whether the Sun is up or not. Thermal control runs continuously and gets harder at night rather than easier. The laboratory can be dimmed. Only two lines are genuinely flexible, they are the two longest bars on the chart, and that is the only reason the architecture closes at all.

> **A sheddable load** is one the base can drop for hours without harming the crew or losing the mission. Shedding is how a power system survives a fault. The trap is that the loads easiest to shed are usually the ones the base exists to run: shut the oxygen plant down often enough and the outpost goes back to being a shelter with a delivery bill.

5.2 The night here is not a fortnight

Everyone knows the number: the lunar night lasts about fourteen and three quarter days. It is true at the equator, and it is the reason most people assume a Moon base must be nuclear. It is not true at this site, and a study book should say so plainly rather than repeat a figure because it is famous.

The base stands on the rim of Shackleton crater, near the south pole. There the Sun does not rise and set. It circles the horizon, never climbing more than about 1.5 degrees above it, because the Moon's spin axis is tilted only 1.54 degrees to the plane of its orbit around the Sun. Darkness at such a place is not the Sun going down. It is a ridge or a crater wall getting in the way, which is why the array stands upright like a fence rather than lying flat like a roof, and why it turns about one axis rather than two.

That makes the length of the night a question about terrain, and the answer depends on how finely the ground has been mapped. Table 1.1 gave the figures: 240 days of continuous light with a longest gap of about 1.5 days on the 240 metre map, 65 to 66 hours at the best rim clusters and 112 hours on Connecting Ridge on the 20 metre map. The finer the map, the longer the worst gap, because a finer map sees smaller obstacles. Nobody has yet computed the shadow history of the patch of ground these modules will stand on, so the honest range for the design is 1.5 to 5 days, and a designer sizes for the longer end of it.

5.3 Could a battery carry the night?

Here is a piece of arithmetic that settles an argument people have been having since the 1960s, and it takes four lines. Work an idea to the point where it fails before spending a year on it.

WORKED EXAMPLE 5.1 · COULD A BATTERY CARRY THE EQUATORIAL NIGHT?

A reduced habitat load of 27.6 kilowatts, a night of 14.7 days, and lithium ion cells at a generous 250 watt hours per kilogram.

1. Turn the night into hours.
   14.7 d × 24 h/d = 352.8 h
2. Multiply by the load to get the energy the base has to store.
   27.6 kW × 352.8 h = 9 737 kWh, call it 9 740 kWh
3. Divide by the specific energy of the cells. Watch the units: 250 Wh/kg is 0.25 kWh/kg.
   9 740 kWh ÷ 0.25 kWh/kg = 38 960 kg
4. State it in tonnes, and price it at the landed rate.
   38 960 kg = 39 t, near 39 billion dollars

ANSWER About 39 tonnes of bare cell, and that is the optimistic figure. It ignores the casing, the wiring and the thermal control, and it assumes the cells can be run flat. Allow the usual 80 per cent depth of discharge and the figure passes 48 tonnes. The whole base lands 48.3 tonnes. A battery that carries the equatorial night therefore weighs as much as the entire outpost, which is how you know the idea is dead.

WHERE PEOPLE GO WRONG

It is tempting to conclude from that arithmetic that a Moon base must be nuclear. At the equator, yes. At this site the honest answer is weaker and more interesting. A 66 hour spell at 27.6 kilowatts needs about 1,800 kilowatt hours, which is roughly 7.3 tonnes of cell. The design sets that against a six tonne reactor and calls the battery the heavier of the two. Six tonnes is not what a reactor weighs. It is the goal NASA set for the class, and the point design NASA drew against that goal came out at 10,046 kilograms across the three units it lands, which the study says in as many words. Put ten tonnes into the comparison and the battery is the lighter one.

5.4 Four ways to hold the base up in the dark

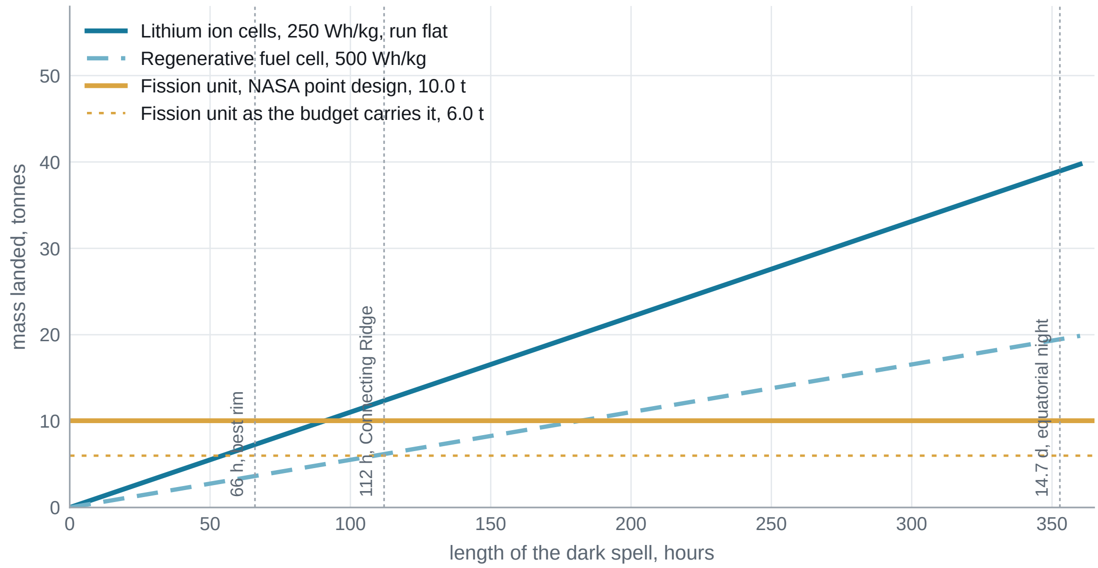

Figure 5.1. How much each way of surviving the dark weighs, against how long the dark lasts, at the 27.6 kilowatt habitat load. The two stores grow with the length of the spell; the reactor does not. The three dotted lines are the dark spells that matter: the best rim site on the fine map, Connecting Ridge on the same map, and the equatorial night that a base anywhere else would face. Read where the lines cross and the argument of this chapter is on the page.

Table 5.1. Options for carrying the base through darkness.

Option	Mass for a 14.7 day night	Mass for a 66 hour spell	Verdict
Lithium ion cells only 250 Wh/kg of bare cell, run flat	39 t	7.3 t	Lightest for a short spell, hopeless for a long one
Regenerative fuel cell hydrogen and oxygen, 500 Wh/kg	19 t	3.6 t	Lighter still, but adds gas storage and still follows the Sun
Fission surface power 40 kWe class unit	10.0 t	10.0 t	Independent of the horizon, and the only option that scales
Fission plus a 210 kWh buffer **the architecture chosen here**	**11.2 to 11.8 t**	**11.2 to 11.8 t**	**Buffer covers reactor transients and peak load**

Landed cost at about a million dollars a kilogram makes each tonne roughly a billion dollars. The fission mass is NASA's own point design, 10,046 kilograms for the three units it lands, and not the 6,000 kilogram goal the budget carries. The architecture row adds a packaged 210 kilowatt hour battery. Only the reactor rows hold their figure across both columns, and that is the whole argument of this chapter in one column of a table.

Read across the reactor rows and the argument has changed shape. At ten tonnes the reactor is not the lightest way to cross a 66 hour shadow. The cells are lighter, at 7.3 tonnes, and the fuel cell lighter again at 3.6. So the reactor has to win on the two things a battery cannot do at all. It does not care where the local horizon is, which matters when nobody has measured the shadow over the ground the modules will stand on. And it scales: the unit that holds the habitat up also runs the oxygen plant straight through the dark, while a store large enough to do that has to be bought over again every time the load grows.

There is one honest caveat, and the design document states it. No 40 kilowatt reactor for the lunar surface exists. The nearest thing ever run is KRUSTY, a 1 kilowatt class test reactor that NASA and the Department of Energy operated in Nevada in March 2018, so the step is from one kilowatt to forty. If the unit is late the base does not stop, because the array and the store carry the pressurised modules while the Sun is up. The oxygen plant is what stops, and an outpost that cannot make its own oxygen is a shelter with a resupply bill.

> **10,046 kg** What NASA's point design for a 40 kilowatt lunar reactor came to, against the 6,000 kilogram goal it was drawn to meet

QUESTIONS

1. A dark spell at this site lasts 66 hours and the reduced habitat load is 27.6 kilowatts. Find the energy a battery would have to hold and the mass of cell at 250 watt hours per kilogram, then set it against the reactor. APPLY
2. Explain, in terms of the angle at which light strikes a surface, why a vertical array beats a horizontal one at this site and would lose badly at the equator. EXPLAIN
3. Two illumination studies of the same ridge disagree: one gives the longest dark spell as 1.5 days, the other as 66 hours. Suggest why the finer study gives the longer answer, and say which figure a designer should size for. EXTEND
4. Figure 5.1 shows the cell line crossing the reactor line at about 91 hours. Show the arithmetic behind that crossing point, and say what happens to it if the reactor is the 6 tonne goal rather than the 10 tonne design. CALCULATE
5. The greenhouse is the largest sheddable load. Describe what actually happens to the plants if it is shed for 66 hours, and what that means for how often the base can afford to do it. EXTEND
6. KRUSTY produced about one kilowatt. List three things that get harder, and one that gets easier, when a reactor is scaled from one kilowatt to forty. EXPLAIN

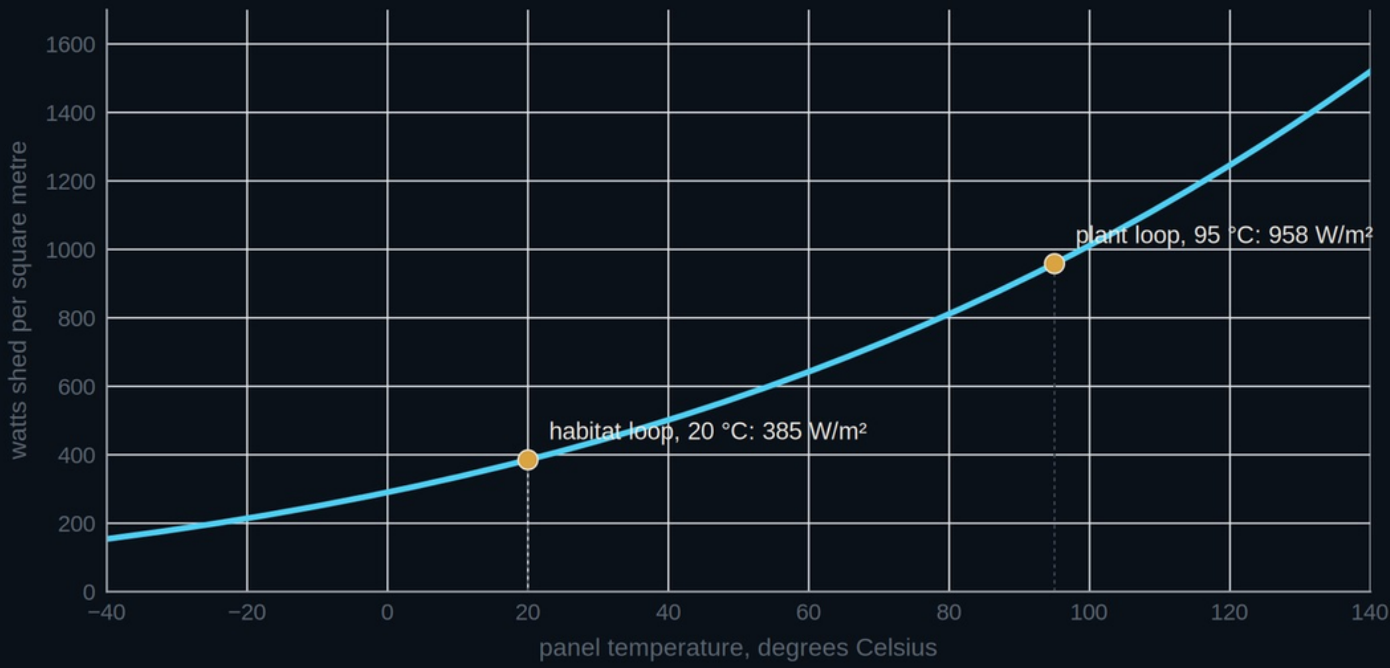

Thermal control

Every watt this base makes has to leave as light, because there is nothing to carry it and nowhere to carry it to. A square metre of radiator at 20 degrees Celsius sheds 385 watts, and the base has 46 kilowatts to lose.

WHAT THIS CHAPTER WILL TEACH YOU

1. Say why vacuum leaves radiation as the only way out, and what that costs in panel area.
2. Use the Stefan Boltzmann law to size a radiator, and show why a hot loop needs less of one.
3. Read a heat budget and name the sources that actually set the radiator area.
4. Test the design's eleven hour buffer claim against its own three numbers.

46 kW
PEAK HEAT TO REJECT

118 m²
RADIATOR INSTALLED

6.2 kW
SURVIVAL HEATING IN THE DARK

5.0 hours
BUFFER HOLD WITH NO POWER

6.1 There is only one way out, and it is slow

On Earth a building loses most of its waste heat to the air around it. Warm air rises, cold air takes its place, and the problem solves itself for nothing. There is no wind on the Moon, no water and no air.

So there is no convection and nothing to conduct into, and the only way a watt can leave this base is as radiation. That would be fine if radiation were quick. It is not. A panel at 20 degrees Celsius with an emissivity of 0.92, facing a sky at 3 kelvin, sheds about 385 watts for every square metre of its face, and the base has 46 kilowatts to lose.

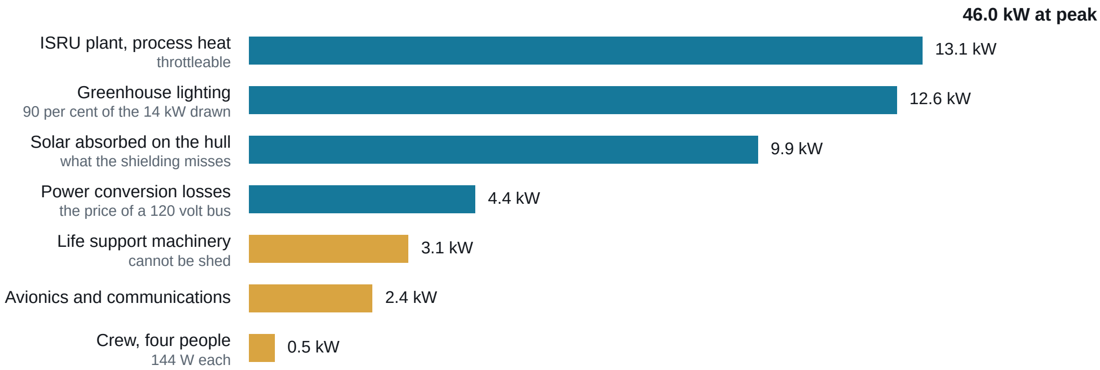

Figure 6.1. Where the heat comes from at peak, in kilowatts. The plant and the grow lights are 25.7 kilowatts between them and both can be throttled. The design calls that two thirds of the load; 25.7 over 46.0 is 56 per cent, so shedding both still leaves 20.3 kilowatts to reject rather than 15.3. The amber bars cannot be turned off, and the crew are the shortest of them.

Figure 6.1 says where that heat comes from. The oxygen plant, the in situ resource utilisation plant, is the largest single source at 13.1 kilowatts of process heat. The greenhouse is next, and it is the surprise: its lights draw 14 kilowatts and 12.6 of that comes straight back out as heat in the bay, because horticultural diodes turn 42 to 93 per cent of their electricity into photons and the plants keep very little of what lands on them. Sunlight the shielding does not stop adds 9.9 kilowatts.

Then the Sun drops behind the rim and the problem runs the other way. In a dark spell, and through a polar winter in which the ridge itself sits between 50 and 70 kelvin, every square metre that was a radiator becomes a leak. The design puts the survival load at 6.2 kilowatts, spent keeping the fluid lines above freezing and the crew volume habitable at the moment the array is making nothing. That 6.2 has no published counterpart and nothing in the design to check it against, which matters, because every claim made for the buffer is that number divided into.

WHERE PEOPLE GO WRONG

Space is cold, so cooling in space must be easy. The temperature of space is not the difficulty. The connection to it is. A vacuum flask keeps coffee hot for hours in a room at 20 degrees Celsius, and it does that by removing convection and conduction and leaving only radiation. This base is a vacuum flask with four people and 46 kilowatts inside it. The 3 kelvin sky is a perfect sink and it is very nearly useless, because the rate at which heat leaves is set by the temperature of the panel and not by the temperature of the sky.

6.2 The fourth power is the whole argument

A surface radiates not in proportion to its temperature but to the fourth power of it, so a modest lift in temperature buys a large gain in output. That one fact decides how the base is plumbed. The habitat loop runs near 20 degrees Celsius, because that is where crew and equipment have to be held. The plant's loop runs at 95, and the difference between those two numbers is worth twenty square metres of panel.

THE ONLY FORMULA IN THIS CHAPTER

$$Q = \varepsilon \sigma A T^4$$

Q is the heat radiated away in watts. ε is emissivity, from 0 to 1; the panels here are 0.92. σ is the Stefan Boltzmann constant, 5.67×10^{-8} watts per square metre per kelvin to the fourth. A is the radiating area in square metres. T is the absolute temperature of the panel in kelvin: add 273.15 to a Celsius reading. This is what a panel emits, not what it nets; subtract what it takes back in from the ground and the low Sun. σ is fixed by nature and T is fixed by the process being cooled, so A is the only term on the right that a designer buys.

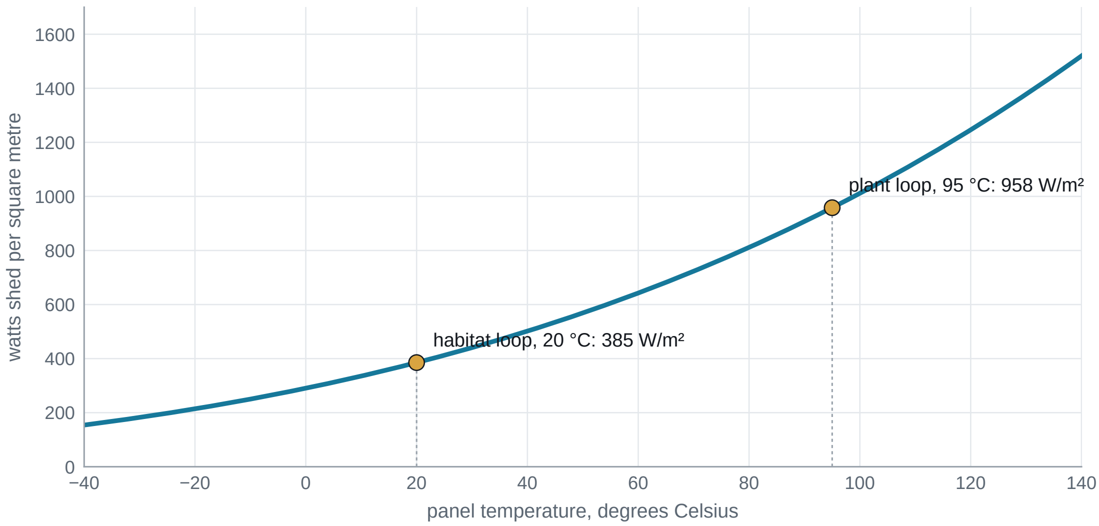

Figure 6.2. What a square metre of panel sheds to a 3 kelvin sky, against the panel's temperature, for an emissivity of 0.92. The curve is the fourth power made visible. The two marked points are the base's two loops, and the vertical distance between them is why the plant runs its own hot loop rather than dumping into the habitat water.

WORKED EXAMPLE 6.1 · WHAT A HOT RADIATOR BUYS YOU

A panel of emissivity 0.92 facing a 3 kelvin sky, first at the habitat loop temperature of 20 degrees Celsius and then at the 95 degrees the plant loop runs at. The plant rejects 13.1 kilowatts.

1. Put both temperatures into kelvin, because the law does not work in Celsius.
   $20\ °C + 273.15 = 293.15\ K$ $95\ °C + 273.15 = 368.15\ K$
2. Raise each to the fourth power.
   $293.15^4 = 7.385 \times 10^9\ K^4$ $368.15^4 = 1.837 \times 10^{10}\ K^4$
3. Multiply each by ε and by σ to get the watts off every square metre.
   $0.92 \times 5.67 \times 10^{-8} \times 7.385 \times 10^9 = 385\ W/m^2$, and the same working at 368.15 K gives $958\ W/m^2$
4. Divide the plant's load by each figure to get the area it would need.
   $13\,100\ W \div 385\ W/m^2 = 34.0\ m^2$ $13\,100\ W \div 958\ W/m^2 = 13.7\ m^2$

ANSWER A lift of 75 degrees takes the panel from 34.0 square metres to 13.7, twenty square metres less for the same 13.1 kilowatts. That is why the plant does not dump into the habitat water at 20 degrees but runs a loop of its own at 95. The design installs 28 square metres rather than 13.7, and the surplus is margin for dust on the face and for the sunlight the panel absorbs. Rejecting heat at the highest temperature a process will tolerate is the cheapest move available in a thermal design, and it costs no mass at all.

No ammonia inside, ever. Ammonia stays liquid across the whole range the outside loop sees; water in an external line would freeze in the first dark spell. But a leak of ammonia into a crew module is fatal, so it never crosses the pressure boundary. The two fluids meet in an interface heat exchanger inside a connector, beyond the habitable volume, and nowhere else.

6.3 Ten wings, and five hours

Table 6.1. Radiators, module by module.

Element	Area	Rejects	Loop	Note
Habitation core, two wings	34 m²	13.2 kW	20 °C	Crew, avionics and galley
Greenhouse bay, two wings	32 m²	12.6 kW	22 °C	Matched to the diode load
Laboratory, two wings	14 m²	5.4 kW	18 °C	Includes the freezer reject
Airlock and cupola, one wing each	10 m²	3.7 kW	20 °C	Peaks during suit servicing
ISRU plant, standalone loop	28 m²	13.1 kW	95 °C	High grade heat, so a smaller panel does more
Installed	**118 m²**	**48.0 kW**		**4 per cent margin on the 46 kW peak**

Each area is worked at the temperature of that module's own loop and not at one figure for all of them. A clean face seeing nothing but sky sheds 385 watts a square metre at 20 degrees Celsius, 396 at the greenhouse's 22 and 958 at the plant's 95, because the rate climbs as the fourth power of the absolute temperature.

Ten wings on five modules, rather than one radiator farm beside them. Sizing each module its own panels costs about 15 per cent more area than pooling the load would, and the 15 is the design's own estimate. What it buys is isolation: the worst single puncture, a core wing, costs about 14 per cent of capacity rather than all of it.

Between the loops and the radiators sits 620 kilograms of Glauber's salt, sodium sulfate decahydrate, which melts at 32.4 degrees Celsius. Melting takes in energy without the salt getting any hotter, so the buffer flattens the peaks and lets the panels be sized for the average rather than the maximum. The design then says the buffer holds the outside lines liquid for eleven hours with no power at all, and calls that the single number the night architecture is built around. Do the division. 620 kilograms at 180 kilojoules a kilogram is 112 megajoules, and 112 megajoules at 6.2 kilowatts is 18,000 seconds, which is five hours. Eleven would need either a draw of 2.8 kilowatts, meaning the buffer keeps the plumbing liquid and lets the crew volume cool, or 1,364 kilograms of salt, twice what the design lands. This book prints five hours, because that is what the design's own three numbers give, and a frozen line cannot be recovered in the field.

Dust attacks a radiator from one end only: it raises how much sunlight a panel takes in rather than lowering how much heat it gives out. The design's table drops the emissivity to 0.68 after two years; NASA's measurement of a fully dusted surface gives about 0.83, and this book uses the measured figure.

> **5.0 hours** What 620 kilograms of salt at 180 kilojoules a kilogram actually holds against a 6.2 kilowatt survival load, against the eleven hours the design claims

QUESTIONS

1. The buffer is 620 kilograms of salt hydrate at an assumed latent heat of 180 kilojoules per kilogram. Find its capacity, then find how long it carries the full 6.2 kilowatt survival load. The design claims eleven hours: find the two ways that could be true. APPLY
2. A radiator wants a low absorptivity and a high emissivity. Say which of the two dust actually spoils, and use that to explain why a vertical panel suffers far less than a horizontal one. EXPLAIN
3. This chapter carries the power conversion losses at 4.4 kilowatts and the power chapter carries 1.1 kilowatts for the same bus. Work out what the disagreement is worth in panel area, compare it with the margin the radiator table holds, and say which figure you would carry. EXTEND
4. Explain why the greenhouse, which draws 14 kilowatts of electricity, gives back 12.6 kilowatts as heat, and where the other 1.4 kilowatts goes. EXPLAIN
5. The plant's 28 square metres would shed 26.8 kilowatts if it were clean and saw nothing but sky. It rejects 13.1. Account for the difference. CALCULATE
6. Suppose the ISRU plant were redesigned to run its loop at 150 degrees Celsius. Find the panel area it would need for 13.1 kilowatts, and name one thing that would get harder. CALCULATE

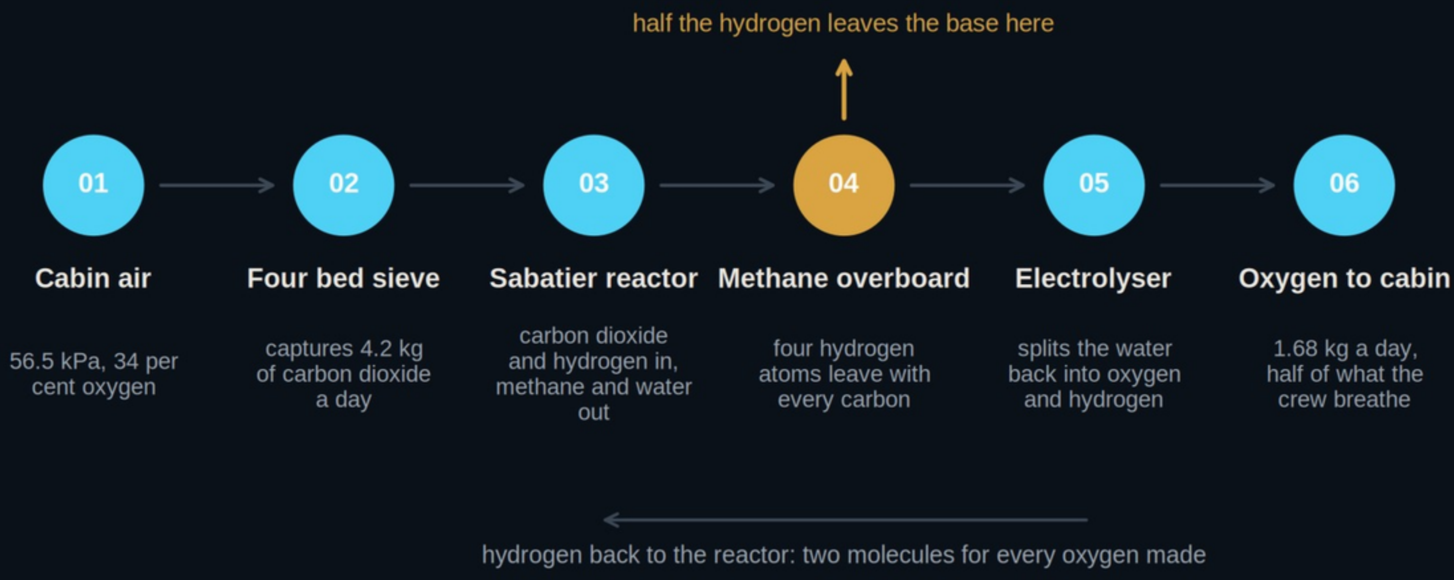

Air and water

The water loop hands back 98 per cent of what it is given. The air loop hands back about half, and the reason is not that anybody built it badly.

WHAT THIS CHAPTER WILL TEACH YOU

1. Build the daily consumable ledger for a crew of four, and say which lines close and which do not.
2. Use the stoichiometry of the Sabatier reaction to show why oxygen recovery stalls at one half.
3. Tell the fraction of carbon dioxide reduced from the fraction of oxygen returned, and say why the two differ.
4. Judge whether a base should ship hydrogen from Earth or make its own oxygen, and on what grounds.

98 per cent
WATER RECOVERED

about 50 per cent
OXYGEN LOOP CLOSED

1.65 kg
OXYGEN SHORT EACH DAY

2 014 kg
LIFE SUPPORT MASS

7.1 What four people need

Land every kilogram of air, water and food that four people need and it comes to 7.5 tonnes a year. The whole base lands 48.3 tonnes, once, and that has to cover the reactor, the excavator and everything else. Life support is not a comfort system. It is the difference between an outpost and a campsite with a return ticket.

The daily bill is 20.64 kilograms for the four of them, 5.16 kilograms each. Most of it is water: 8.00 kilograms to drink and to cook with, another 6.40 to wash and to flush. Then 3.36 kilograms of oxygen, 2.72 of dry food, and 0.16 of nitrogen to replace what leaks out through the hull. The loops hand 15.85 kilograms of that back. The other 4.79 has to arrive from somewhere.

The three large lines close to very different degrees. Water is nearly shut: 98 per cent recovered, 0.29 kilograms a day still outstanding across both streams. Food is barely shut at all, because the greenhouse returns 0.03 kilograms of dry matter against a demand of 2.72. Air sits between them, and air is the interesting one. A little over half the oxygen the crew breathe comes back, and it stops there for a reason that has nothing to do with how well the machinery was built. What is left over is 1.65 kilograms of oxygen a day, 297 kilograms a rotation, for as long as anybody lives here.

> **The exploration atmosphere.** The cabin runs at 56.5 kilopascals, 8.2 pounds per square inch, with 34 per cent oxygen. Low pressure with rich air leaves a crew member less dissolved nitrogen to purge before going outside, so the prebreathe is short. Oxygen partial pressure is 19.2 kilopascals against 21.2 at sea level, so the crew live permanently a little short of breath.

7.2 Why the oxygen loop stops at one half

A four bed molecular sieve takes 4.2 kilograms of carbon dioxide a day out of the cabin air and regenerates its beds by opening them to vacuum, which is free on the Moon and is not on a station. The gas goes to a Sabatier reactor, meets hydrogen over a hot catalyst and leaves as methane and water. The water is split. The oxygen goes to the crew, the hydrogen goes back to the reactor, and the methane goes overboard. Nothing in that description is wrong. It is also not enough.

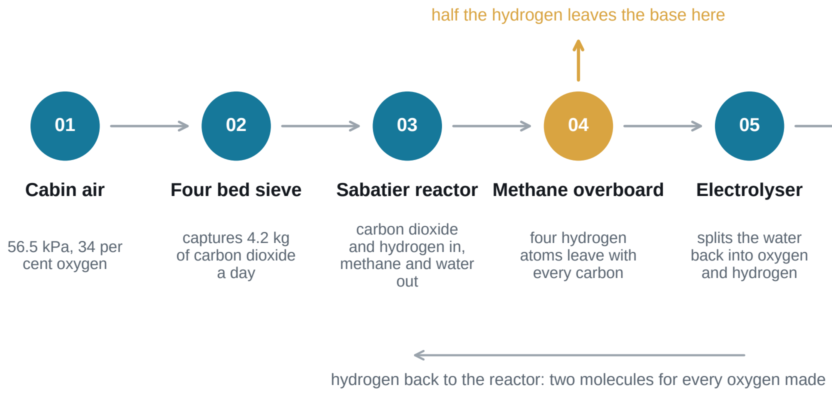

Figure 7.1. The air loop, drawn as a line and not as a ring, because it does not close. Everything the crew exhale is captured. At step 04 half of the hydrogen the reactor was fed walks out of the base inside the methane, so step 06 returns 1.68 kilograms of oxygen a day against the 3.36 the crew breathe. Of the difference, 0.03 kilograms comes from the greenhouse and the remaining 1.65 has to be made on the surface.

Count the atoms. The reaction is $CO_2 + 4H_2 \rightarrow CH_4 + 2H_2O$, and every molecule of methane thrown overboard takes four hydrogen atoms with it. The electrolyser hands back two molecules of hydrogen for each molecule of oxygen it makes. The reactor spends four on each molecule of carbon dioxide it reduces. Two over four is one half, and no catalyst, temperature or pressure alters it, because the quantities being divided are atoms in a molecule.

The hardware itself has flown. A Sabatier assembly worked on the International Space Station from 2010 until 2017, when its catalyst was found poisoned. The point is not that the machine is unproven; it is that a perfect machine would still return half.

WORKED EXAMPLE 7.1 · TWO FRACTIONS THAT ARE EASY TO CONFUSE

Four crew breathe 3.36 kilograms of oxygen a day and exhale 4.2 kilograms of carbon dioxide. Molar masses: oxygen 32 grams per mole, carbon dioxide 44. The reaction is $CO_2 + 4H_2 \rightarrow CH_4 + 2H_2O$.

1. Count the oxygen in moles, because a reaction counts molecules and not kilograms.
   3 360 g ÷ 32 g/mol = 105 mol a day
2. The electrolyser frees two molecules of hydrogen for each molecule of oxygen it makes.
   2 × 105 mol = 210 mol of hydrogen a day
3. The reactor spends four of them on each carbon dioxide it reduces, so the hydrogen reaches this much carbon dioxide.
   210 mol ÷ 4 = 52.5 mol of carbon dioxide
4. Set that against the carbon dioxide the crew actually exhale.
   4 200 g ÷ 44 g/mol = 95.5 mol; 52.5 ÷ 95.5 = 55 per cent reduced
5. Each carbon dioxide reduced gives back one oxygen molecule. Set that against the oxygen breathed.
   52.5 mol × 32 g/mol = 1 680 g = 1.68 kg; 1.68 ÷ 3.36 = 50 per cent returned

ANSWER The reactor reaches 55 per cent of the carbon dioxide and returns 50 per cent of the oxygen. The two fractions are not the same number and it is worth seeing why, because almost everybody who quotes one of them quotes the other by mistake. The crew breathe in 105 moles of oxygen and breathe out only 95.5 moles of carbon dioxide, since some of the oxygen leaves the body in water rather than in carbon dioxide. The reactor is working on the smaller stream, so the same 52.5 moles is a larger slice of the carbon dioxide than of the oxygen.

WHERE PEOPLE GO WRONG

The 50 per cent gets read as an engineering shortfall, as though a hotter reactor or a cleverer catalyst would close it. Look at the reaction again. The only quantities in it are the number of hydrogen atoms in a methane molecule and the number in a water molecule. A reactor converting every molecule it is fed still loses half the hydrogen, because half of it is inside the methane. There are three ways round this and none of them is a better reactor: bring hydrogen in from outside the loop, crack the methane and take its hydrogen back, or make oxygen from something that is not cabin air.

78 kg Lithium hydroxide carried. If the sieve fails the crew fall back on powder that takes carbon dioxide out of the air and never gives it back. Flight hardware needs about 28 kilograms for a 96 hour supply; the design carries 78, a margin of 2.8 times on a backup nobody expects to open.

7.3 The loop that closes, and the gap that gets dug

Water is the success story, and a recent one. The station recovered 93 to 94 per cent of what it collected until a brine processor went up in March 2021 to work on what the urine still leaves behind, and by 2023 the figure was 98 per cent. Five points came out of squeezing brine. The largest stream is not the urine processor at all: it is the cabin air, giving up 5.9 kilograms a day of condensate on cold surfaces, ahead of 4.3 of grey water and 4.1 of distillate.

The oxygen gap will not yield to any of that, so it has to be closed some other way, and there are only three routes out. Ship hydrogen from Earth, 0.38 kilograms a day and 139 kilograms a year, which is nothing against 48.3 tonnes and buys 91 per cent rather than 100. Crack the methane and take its hydrogen back, which is NASA's stated minimum of 75 per cent and has not yet flown. Or make oxygen out of the ground, which costs a tonne of plant and 14 kilowatts and owes Earth nothing. This base takes the third, and carries methane pyrolysis in increment four as a demonstration rather than as something the crew depend on. Mass is not the reason. A base that ships hydrogen is a very efficient camp, and a base that makes its own oxygen is the beginning of an industry.

Table 7.1. Four ways to close the oxygen gap.

Option	Closure	What it costs	Verdict
Sabatier alone flown 2010 to 2017	about 50%	nothing new	The baseline, and the gap has to be filled some other way
Hydrogen shipped from Earth 0.38 kg a day	about 91%	139 kg a year landed	Cheap in mass, and a dependency that never ends
Methane pyrolysis crack the CH_4, keep its hydrogen	75%	a plasma pyrolysis assembly	NASA's stated minimum, and not yet flown
Oxygen from regolith **the architecture chosen here**	**100%**	**1 t of plant, 14 kW**	**Independent of Earth, and paid for in electricity**

The 91 per cent in the second row is a ceiling and not a target. It is set by the crew themselves, who breathe in more molecules of oxygen than they breathe out as carbon dioxide.

QUESTIONS

1. Shipping hydrogen is said to close the oxygen loop to about 90 per cent and not to 100. The crew breathe 3.36 kilograms of oxygen a day and exhale 4.2 kilograms of carbon dioxide. Use moles to find the ceiling, and say where the missing oxygen goes. APPLY

2. Explain, without using the word efficiency, why water recovery reaches 98 per cent while oxygen recovery stalls at 50 per cent. EXPLAIN

3. The water loop counts 2.1 kilograms a day of greenhouse transpiration as recovered, and no line of the demand column pays for watering the plants. State the assumption this hides, and say what you would need to know to test it. EXTEND

4. Find the mass of hydrogen that would have to be shipped each day to reduce every molecule of carbon dioxide the crew exhale, and show that it comes to about 0.38 kilograms. CALCULATE

5. Methane pyrolysis returns hydrogen from the methane. If it recovered every hydrogen atom, what closure would the loop reach, and why is NASA's stated figure lower? EXPLAIN

6. The lithium hydroxide backup lasts 96 hours at flight scale. Work out how long 78 kilograms lasts four people, and say what the base would have to do in that time. CALCULATE

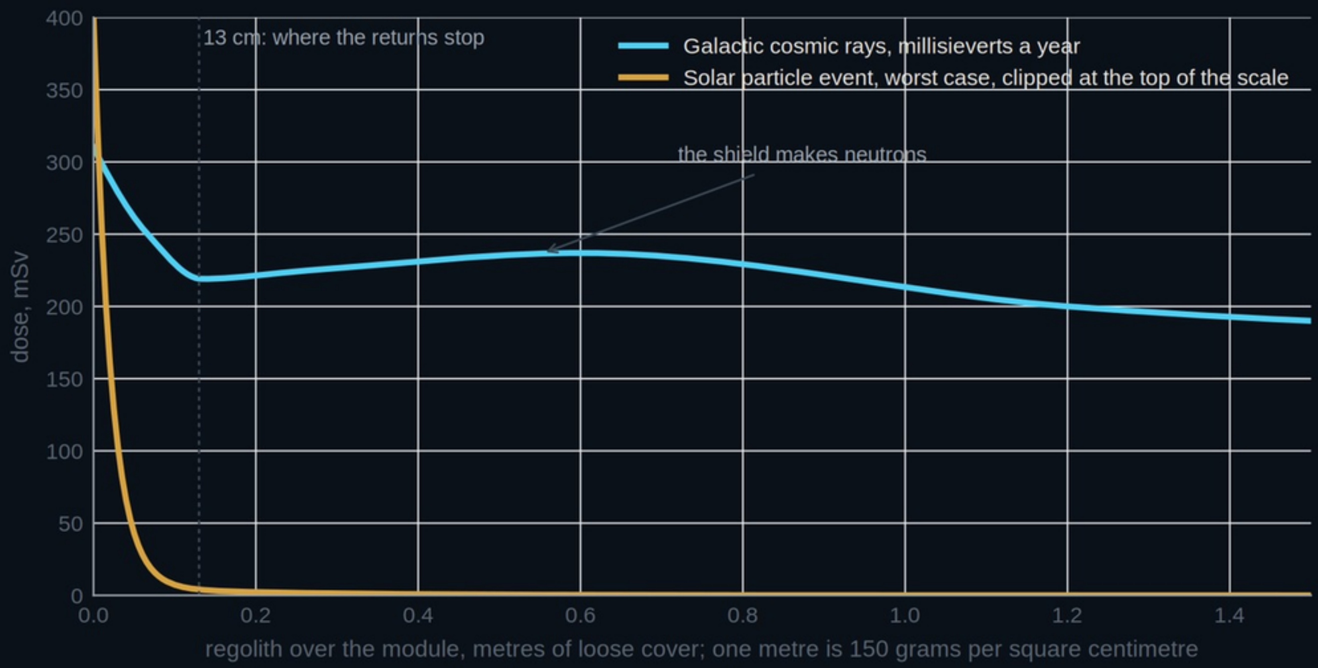

Radiation

Nothing here stops a particle except the mass put in its way, and this base lands 48.3 tonnes in total. The cover that brings the routine dose down has to be dug out of the ground the base is standing on.

WHAT THIS CHAPTER WILL TEACH YOU

1. Tell a galactic cosmic ray from a solar particle event, and say why only one of them can be stopped.
2. Work in areal density rather than in thickness, and turn a depth of regolith into grams per square centimetre.
3. Keep the gray and the sievert apart, and say which question each of them answers.
4. Say where the shielding curve stops paying, and what it is that stops it.

380 mSv a year
ROUTINE DOSE, UNSHIELDED

600 mSv
CAREER LIMIT

13 cm of regolith
WHERE THE RETURNS STOP

30 minutes
WARNING OF AN EVENT

8.1 Two threats, and only one of them stops

Three centimetres of soil will stop the worst solar storm on record. Two metres will not stop the ordinary background, and no depth anybody could dig ever will. Those two sentences are most of this subject, and they run the opposite way to what everybody expects.

The two threats have almost nothing in common. Galactic cosmic rays are a thin, continuous rain of very high energy nuclei from beyond the solar system, typically a few gigaelectronvolts, present every second of every day. Solar particle events are the opposite in every way: floods of protons thrown off the Sun by a flare, mostly 10 to 500 megaelectronvolts, lasting hours to days. About one a year is worth taking cover from.

So the base has two answers. Against solar protons, a storm shelter under the habitation core deck, ringed by water and waste tanks the base lands anyway: 12 grams per square centimetre. Behind that the worst event on record leaves a few tens of milligray equivalent in the blood forming organs, against a limit of 250 in any thirty days, and nobody works further than twenty minutes from an airlock, because the first protons arrive half an hour after the flare. Against galactic cosmic rays, two metres of regolith. That cover is not built at landing: the modules stand on their skirts under an open sky, at about 380 millisieverts a year, until the excavator has moved it.

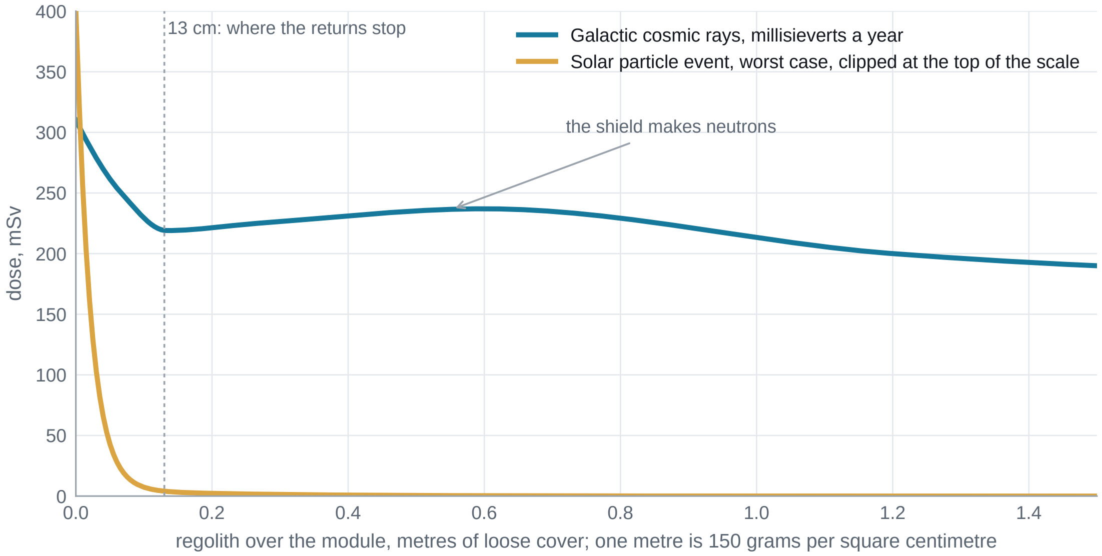

Figure 8.1. Dose against the depth of loose regolith, drawn through the points Matthiä and Berger publish. The two curves are different quantities sharing one axis: the galactic one is a rate, in millisieverts a year, starting at 310, the dose calculated for a body rather than the 380 field estimate the design carries; the solar one is a single worst case event, clipped at the top of the scale. Watch where each of them stops falling. The solar curve has lost five sixths of its dose in the first three centimetres and nearly all of it by thirteen. The galactic curve has taken its whole cut by thirteen as well, and then turns back up while the shield makes neutrons. The design's two metres is off the right hand edge of this chart.

8.2 Weighed, not measured

Shielding is not measured in thickness. It is weighed. What stops a particle is the number of nuclei standing in its path, not the distance it travels, so the quantity is areal density: the mass in the path, per unit of area. Two shields of the same areal density are nearly the same shield whatever they are made of.

THE ONLY FORMULA IN THIS CHAPTER

$$\sigma = \rho \times d$$

σ is areal density, the mass in the path per unit of area, in grams per square centimetre. ρ is bulk density in grams per cubic centimetre: loose regolith is about 1.5, compacted about 1.8, and the top of the scale for packed soil is about 2.0. d is depth in centimetres. Because σ is what matters, a metre of loose soil at 150 grams per square centimetre and 83 centimetres of compacted soil are the same shield.

Two units, and the difference is not pedantry. The gray is absorbed dose: one joule of energy left in each kilogram of matter, with no biology in it at all. The sievert is dose equivalent, that same energy weighted for the harm the radiation does, so a gray of heavy nuclei is worth several sieverts and a gray of X rays is worth one. Now watch a shield work. The first 20 grams per square centimetre take the dose equivalent from 0.84 to 0.60 millisieverts a day, a cut of about a quarter, while the absorbed dose goes from 0.25 to 0.28 milligrays and does not fall. Measured in grays the shield achieves nothing. It is not stopping the energy; it is changing what kind arrives.

Short term organ limits use a third quantity, the milligray equivalent, weighted for prompt injury rather than cancer risk. The career limit is 600 millisieverts of effective dose, one figure for every age and both sexes since 2022, a decision about risk rather than a physical threshold. Only one reading has ever been taken on the surface: Chang'e 4's 1,369 microsieverts a day in 2019, charged particles only, in a small detector, which works out to about 500 millisieverts a year and sits above the 380 the design carries.

WORKED EXAMPLE 8.1 · WHAT TWO METRES OF COVER ACTUALLY WEIGHS

Loose regolith at 1.5 grams per cubic centimetre, a cover 2.00 metres deep, a habitation core 8.0 metres in diameter, and a base that lands 48.3 tonnes in total.

1. Put the depth into centimetres, because the density is per cubic centimetre.
   $d = 2.00\text{ m} = 200\text{ cm}$
2. Multiply density by depth for the areal density.
   $1.5\text{ g/cm}^3 \times 200\text{ cm} = 300\text{ g/cm}^2$
3. Turn that into mass over a square metre. One square metre is 10,000 square centimetres.
   $300\text{ g/cm}^2 \times 10\,000\text{ cm}^2/\text{m}^2 = 3\,000\,000\text{ g/m}^2 = 3.0\text{ t/m}^2$

4. Take the flat area under the core, a circle 8.0 metres across, and multiply by that.
   $A = \pi r^2 = 50.3\ m^2$; $50.3\ m^2 \times 3.0\ t/m^2 = 151\ t$

ANSWER About 151 tonnes over one module, and that is a floor rather than an estimate, because a cover has to arch over a dome and reach the ground beyond its footprint. The whole base lands 48.3 tonnes, reactor and excavator included. Spread every gram of that over the core instead and you would have 48,300 kg ÷ 50.3 m² = 960 kg/m², which is 96 grams per square centimetre, under a third of the target. Shielding at this scale cannot be launched. It has to be dug, which is why the excavator is a life support system with a bucket on it.

WHERE PEOPLE GO WRONG

The dentist puts a lead apron on you, so lead must be the best shield there is. For a dental X ray it is, because those photons are low in energy and a heavy nucleus soaks them up. A galactic cosmic ray is not a photon. It is a nucleus at a few GeV, and when it strikes a lead nucleus it smashes it, and the fragments go on as secondaries. Per kilogram the ranking runs the other way and what it runs on is hydrogen content: liquid hydrogen first, then polyethylene at 14.4 per cent hydrogen by mass, then water at 11.2 per cent, then regolith, then aluminium. Regolith is used anyway, because 300 grams per square centimetre of a mediocre shield beats 20 of an excellent one that had to be launched, and against a solar particle event 20 of anything already finishes the job.

8.3 Where the curve stops paying

Draw the dose against the depth of cover and the curve does something nobody expects. It falls quickly, and it has finished falling by about 13 centimetres. Matthiä and Berger put the dose equivalent at 0.84 millisieverts a day on the open surface and 0.60 behind 20 grams per square centimetre, then have it rise again to 0.65 at 90. NASA's own standard is blunter: between 40 and 100, more shielding has a negligible or a negative effect. The design's chart used to put the knee at 250, more than an order of magnitude too deep, which made two metres look like the depth where value runs out. It ran out at a hand's breadth.

Two mechanisms are at work here, easy to swap over. Fragmentation is why the first ten centimetres help: a heavy nucleus striking a nucleus in the shield shatters, and the light fragments lay down a far thinner track of ionisation than an iron nucleus does, so the average quality factor falls from about 3.4 to 2.3. What stalls the curve is not fragmentation but neutrons. Past about 20 grams per square centimetre the shield makes them faster than it breaks up primaries, and behind 180 they are about half the dose equivalent. The shield is a neutron factory, and it takes a metre of soil before the factory stops growing.

Everything between 13 centimetres and two metres buys back neutrons the shield itself made, and buys them slowly. Bare regolith at 180 grams per square centimetre gives about 200 millisieverts a year against 310 unshielded, so the design's own 126 for two metres is optimistic. This book carries the design's dose table below as the design states it, and prints the published band beside it, because the honest position is that no transport calculation has been run for this hull, this cover and this ground.

Table 8.1. The dose budget, in the design case with the cover in place.

Exposure	Rate	Hours a year	Dose a year
Inside, under two metres of regolith	0.0144 mSv/h	6 570	94.6 mSv
Suited work outside	0.043 mSv/h	912	39.2 mSv
Inside the cupola	0.036 mSv/h	365	13.1 mSv
Transit, Earth to Moon	1.2 mSv/day	192	9.6 mSv
In the storm shelter	0.0144 mSv/h	48	0.7 mSv
Per 180 day rotation			**78.6 mSv**

Read the second row against the first. Suited work is a tenth of the crew's time and a quarter of the dose, and an hour outside is worth three hours in. That is the argument for autonomy on this base: an hour a machine spends at the excavation face is an hour a person does not, and a machine has no career limit. Every row assumes the two metres is in place. Until it is, the rate inside is the open surface rate of the second row and a rotation costs about 200 millisieverts rather than 79. The hours are for a year and the total is for a rotation, so the total row is half the column above it. The published band for bare regolith cover runs from about 150 to 200 millisieverts a year rather than the 126 the first row implies.

> **151 t** What two metres of loose regolith weighs over the habitation core alone, against the 48.3 tonnes the whole base lands

QUESTIONS

1 Loose regolith is 1.5 grams per cubic centimetre and compacted regolith about 1.8. Find the depth of each that gives 300 grams per square centimetre, and say which the excavator should aim for and why. APPLY

2 Explain why a shield can reduce the dose in sieverts while leaving the dose in grays unchanged, and what that tells you about what the shield is doing to the radiation. EXPLAIN

3 The design puts two metres at 126 millisieverts a year. Matthiä and Berger find 200 behind 180 grams per square centimetre, and Horst and colleagues reach 100 only with 400 and a polyethylene liner. Turn those two shields into depths of loose regolith, decide which number to carry, and say what you would want measured. EXTEND

4 Chang'e 4 measured 1,369 microsieverts a day. Convert that to a yearly rate and compare it with the 380 the design uses. Give one reason the two might legitimately differ. CALCULATE

5 The first protons of a solar event arrive about thirty minutes after the flare. A crew member walks at about 1 metre a second in a suit. How far from the airlock can work be planned, and what else limits the answer? APPLY

6 Using Table 8.1, work out how many 180 day rotations a crew member could serve before reaching the 600 millisievert career limit, first with the cover in place and then without it. CALCULATE

CHAPTER NINE

dust travels one way along this line, and each amber boundary is meant to stop it

two curtains at 96 per cent leave 0.04 × 0.04 = 0.16 per cent, and boundaries multiply rather than add

Dust

Lunar dust is broken glass that has never been weathered, and it carries a charge. This base does not plan to clean it off the crew. It plans never to let it cross the hull.

WHAT THIS CHAPTER WILL TEACH YOU

1. Say what makes lunar grains sharp, clinging and small enough to breathe, and what each of those costs.
2. Trace how the limit of 0.3 milligrams per cubic metre was arrived at, and why it sits below every analysis.
3. Use a permissible exposure limit to work out how much dust a pressurised volume may hold, and turn that into a capture requirement.
4. Judge a suit port against a traditional airlock, and say what the design gives up as well as what it gains.

20 per cent
OF SOIL FINER THAN 20 MICROMETRES

0.3 mg per m³
CABIN DUST LIMIT

340 hours
SHORTEST SERVICE LIFE

96 per cent
CURTAIN CAPTURE TARGET

9.1 Broken glass that has never been weathered

Gene Cernan told the Apollo 17 debrief that dust was probably the greatest inhibitor to a nominal operation on the Moon. He had been on the surface for just under 75 hours. This base is planned for a decade.

Terrestrial dust has been rolled by wind and water for millions of years, so its grains are worn round. Lunar regolith has been shattered by micrometeorite impact in vacuum and never weathered at all, so its grains stay angular, fractured and often glassy. Under a microscope they look like broken bottles. They also carry a charge, put on them by solar ultraviolet on the sunlit side and by plasma on the night side, which is why they cling to a glove and will not brush off.

Size is the property that reaches the crew rather than the machinery, and three sizes get run together. About 20 per cent of lunar soil is finer than 20 micrometres by mass. The respirable fraction, fine enough to travel past the airways into the alveoli, is finer than 2.5 micrometres and comes to only 1 to 3 per cent of a mature soil. NASA's requirement is written against neither. It limits everything below 10 micrometres, the size that stays hanging in cabin air, to 0.3 milligrams per cubic metre averaged over six months, and it says in as many words that it does not cover polar dust. This base stands at the pole.

> **2 moonwalks** to a leaking suit on Apollo 12. After about eight hours outside, Pete Conrad's suit was losing 0.25 pounds per square inch a minute against an allowable 0.30. It had not leaked at all before the first walk.

9.2 A limit is a decision, not a measurement

That 0.3 is worth following back, because it is the clearest case in this book of a number that is not a measurement. Four analyses were run on the same animal data, all of it from returned Apollo soil. A benchmark dose analysis of the instillation study gave 0.5 to 1.0 milligrams per cubic metre. Benchmark dose analyses of the two inhalation studies gave 0.6 and 0.9. A no observed adverse effect level analysis of the same inhalation data gave 0.4. NASA's own toxicology group weighed all four and recommended 0.5. An external review committee then set the limit at 0.3, and it was accepted in April 2014.

Nothing had been measured in between. The committee was not picking the number the data supported best; it was deciding how much of the uncertainty to carry on the crew's side, and there was plenty to carry. None of the soil the animals breathed was polar. The exposure is episodic across six months rather than a working shift, and the people exposed cannot go home at the end of it. An earlier survey of expert opinion had run from 0.01 to 3 milligrams per cubic metre, a range of three hundred to one. A limit set below every analysis is not a mistake. It is what caution looks like when it is written down.

9.3 The boundary that never lets it in

Every dust measure falls into one of three groups: stop it entering, take it out once inside, or put up with it. Cleaning costs crew time and is never finished. Putting up with it is what the Apollo record is. After two moonwalks and about eight hours outside, the Apollo 12 suits were worn through the outer layer above the boots and into the insulation beneath, the wrist and hose locks had clogged, and Pete Conrad's suit was leaking. On Apollo 17 the cover gloves used on the core drill were worn through after two walks and thrown away, and Harrison Schmitt's visor sunshade was scratched until he could not see out in some directions.

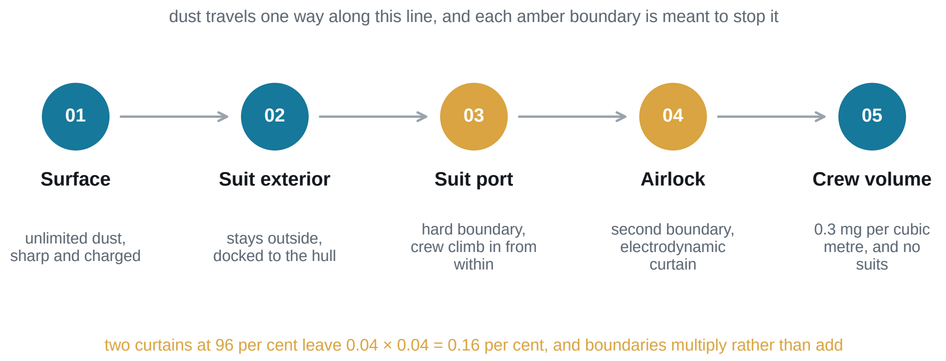

Figure 9.1. The chain of boundaries, drawn as a line because dust travels only one way along it. Only the vestibule behind the hatch is ever cycled, which puts the time before an extravehicular activity under thirty minutes against hours through an airlock. At the last boundary the crew pass an electrodynamic curtain: transparent electrodes driven out of phase, walking charged grains off the edge for 18 watts a face.

WORKED EXAMPLE 9.1 · HOW MUCH DUST MAY THE BASE'S AIR HOLD?

A limit of 0.3 milligrams per cubic metre of cabin air on particles below 10 micrometres, a pressurised volume of 372 cubic metres, and soil of which about 20 per cent by mass is finer than 20 micrometres.

1. Turn the limit into a mass for the whole base by multiplying the concentration by the volume.
   $0.3\ \text{mg/m}^3 \times 372\ \text{m}^3 = 111.6\ \text{mg}$, call it 112 mg

2. Take one kilogram of soil carried in on a suit and find the part the limit counts. Everything below 10 micrometres sits inside the 20 per cent that is below 20, so this is an upper bound.
   $1\ \text{kg} \times 0.20 = 0.20\ \text{kg}$, which is 200 000 mg

3. Compare the kilogram with the allowance.
   $200\,000\ \text{mg} \div 112\ \text{mg} = 1\,786$, about 1 800 times over

4. Turn that round into the fraction the boundaries may pass, then test two curtains at 96 per cent against it.
   $112\ \text{mg} \div 200\,000\ \text{mg} = 0.00056$, that is 0.056 per cent, against $0.04 \times 0.04 = 0.16$ per cent for two curtains

ANSWER One kilogram of soil brought inside carries up to about eighteen hundred times the dust the air may hold, so the boundaries are allowed to pass about 0.06 per cent of it. Two curtains at 96 per cent leave 0.16 per cent, nearly three times the budget. That is why the first boundary is not a filter at all: the suit port keeps the kilogram outside, and the curtain works only on what a glove or a tool carries past it. The true multiple lies somewhere between about ninety and eighteen hundred, because the limit is a six month average and only 1 to 3 per cent of a mature soil is finer than 2.5 micrometres.

BOUNDARIES MULTIPLY, THEY DO NOT ADD

Two stages at 96 per cent do not catch 192 per cent of anything. The fraction that gets through n boundaries in series is p to the power n, where p is the fraction one boundary lets past: 0.04 for a stage that catches 96 per cent. Two stages leave 0.04 × 0.04 = 0.0016, and it is that number, not the 96, that the design has to be sized against.

WHERE PEOPLE GO WRONG

The usual picture is of charged dust levitating off the surface in fountains and settling on everything by itself. The Lunar Atmosphere and Dust Environment Explorer looked for exactly that between 3 and 250 kilometres, set an upper limit of 100 grains per cubic metre, and found instead a thin cloud raised by micrometeoroid impacts. The dust that ruins hardware here is dust somebody moved: a boot, a wheel, a landing. That is the case for boundaries and for sintered roads rather than for cleaning the air. The caveat is that the instrument could not look below 3 kilometres, and nobody has watched dust behave at a site that sits at the terminator for months.

Table 9.1. The trade the suit port makes, measure by measure.

Measure	Traditional airlock	Suit port	Difference
Dust path inside	Every EVA	None	The whole problem
Time before an EVA	Hours	Under 30 minutes	Two EVAs a day, which nobody has yet flown
Atmosphere lost per cycle	The airlock volume	Only the vestibule behind the hatch	Most of the consumable
Prebreathe	Hours from a 14.7 psi cabin	About 20 minutes at 8.2 psi and 34 per cent oxygen	Cuts the largest EVA overhead
Suit stowage inside	Every suit, in the cabin	None, the suits stay docked outside	Volume and dust both stay outside

The gas lost per cycle is asserted rather than sized: it has not been measured for this design. The design carries fifteen minutes of prebreathe; NASA's own chamber run from that same 8.2 psi cabin took twenty, and this book uses twenty. NASA has since moved its testing to a 9.6 psi cabin at 28.5 per cent oxygen, because 34 per cent oxygen is a fire risk, and there the validated prebreathe is 150 minutes. A suit port is also not retrofittable, so the airlock lands with it or never gets it.

> **What the shield has done so far.** Kennedy Space Center flew dust shield panels on the station from 2019, and one landed on Blue Ghost on 2 March 2025 and cleared regolith from glass and radiator samples in the days that followed. That is a two week demonstration and a station exposure. Nobody has yet run one for a decade.

QUESTIONS

1. A curtain catches 96 per cent of the dust reaching it. Find the fraction that gets past one, two and three curtains in series, and say how many are needed to hold it below one part in ten thousand. APPLY

2. No method used to derive a dust exposure limit produced a number below 0.4 milligrams per cubic metre, and the limit was set at 0.3. Explain how that can be the right decision rather than an error, and say what would justify raising it. EXPLAIN

3. The requirement is written against particles below 10 micrometres, and the standard says it does not cover polar dust. The size distribution at the Shackleton rim has never been measured. Say what that does to a 96 per cent capture target, and what a designer should do in the meantime. EXTEND

4. The suit port cuts prebreathe from hours to about twenty minutes. Explain, from the physics of dissolved gas, why a lower cabin pressure with richer oxygen makes that possible, and name the risk it carries. EXPLAIN

5. Pete Conrad's suit leaked 0.25 psi a minute after eight hours outside. Estimate how many hours of EVA a suit on this base would see in a 180 day rotation with two EVAs a day, and say what that implies for how the suits are managed. CALCULATE

6. Table 9.1 says the suit port cannot be retrofitted. Give two reasons that might be so, and say what a designer should decide before the airlock is built. EXTEND

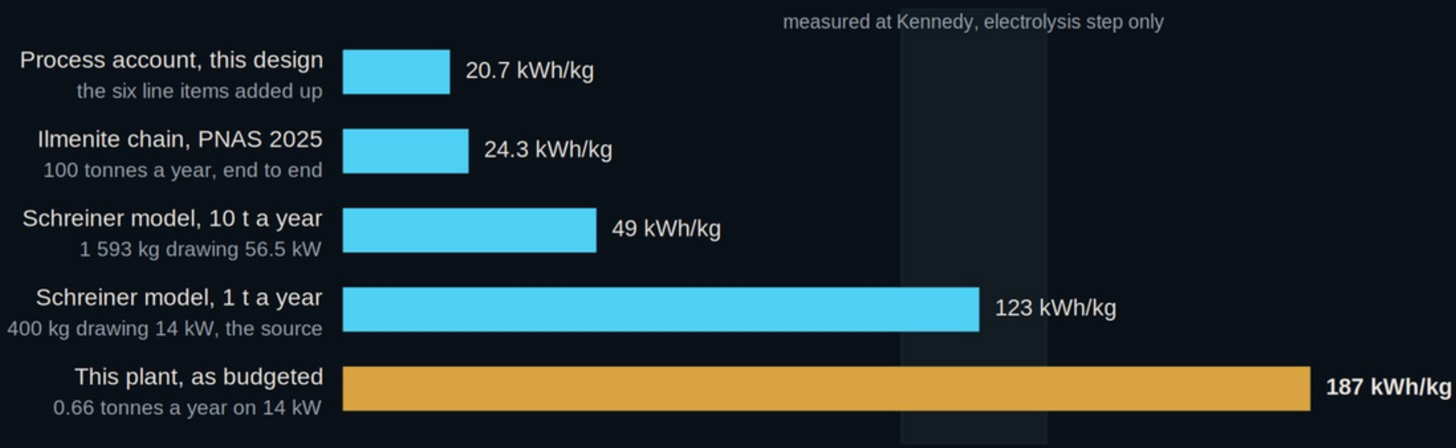

Making it on site

A base that lands 8.9 tonnes of power hardware and then still ships its oxygen from Earth has not made a case for itself. Almost half the mass of the ground under it is oxygen. This chapter is about what it costs to get out, and about one division nobody did.

WHAT THIS CHAPTER WILL TEACH YOU

1. Say why this base mines the ground it stands on rather than the polar ice.
2. Describe molten regolith electrolysis, and what it takes from the soil that hydrogen reduction cannot.
3. Turn a plant's power draw and its daily output into the energy cost of a kilogram of product.
4. Explain why the same chemistry costs 24 kilowatt hours a kilogram in a large plant and 187 in this one.

43 per cent
OXYGEN IN THE GROUND, BY MASS

1.8 kg a day
PLANT OUTPUT

14.0 kW
PLANT DRAW, CONTINUOUS

187 kWh/kg
ENERGY PER KILOGRAM

10.1 The resource is the bulk material

Every other chapter in this part has been about carrying things: 48.3 tonnes landed, 28.7 of it pressure vessel and 8.9 power hardware. Land all that, then meet a freighter twice a year for oxygen, and you have a camp with a delivery bill. The way out is under the crew's boots.

The resource is not buried in the rock; it is the bulk of the rock, and a rover scoops it into the bucket. Samples from eleven sites have come home, the latest from the far side with Chang'e 6 in June 2024, so the composition is known well enough to design against. Mare regolith is basaltic and highland anorthositic, and both are metal oxides almost all the way down. Take the highland column of Table 10.1, multiply each oxide by the fraction of its own mass that is oxygen, add the six lines. Work it out before you read on. The answer is 44.5 per cent, so a tonne of that ground holds 445 kilograms of oxygen, and four crew need 602 in a year.

The design carries 43, which is what to do with a figure that has to survive real soil rather than a table. The highlands are the richer ground of the two, 44.5 per cent against 42.9, and every polar site anybody wants is a highland site.

Table 10.1. What the ground is made of, by mass, from the Lunar Sourcebook. The last column is fixed by chemistry rather than by geology: the mass of oxygen in one formula unit, divided by the mass of the whole unit.

Oxide	Mare	Highland	Oxygen fraction
SiO_2 silica	45.4 %	45.0 %	53 %
Al_2O_3 alumina	14.9 %	24.6 %	47 %
FeO iron oxide	14.1 %	5.7 %	22 %
CaO lime	11.8 %	15.8 %	29 %
MgO magnesia	9.2 %	7.5 %	40 %
TiO_2 titania	3.9 %	0.6 %	40 %
Oxygen, total	**42.9 %**	**44.5 %**	

The total row is the mass weighted sum of the rows above it, computed here rather than quoted. Alumina is 47 per cent oxygen and silica 53 where iron oxide is only 22, so the aluminium rich highlands come out ahead. The design prints 44.6 for mare and 44.3 for highland; neither follows from its own rows, and the pair is the wrong way round.

10.2 Melt it, and pass a current through it

Molten regolith electrolysis is the least clever process in this chapter, which is much of why it was chosen. Heat regolith above 1,600 degrees Celsius and it stops being a powder and becomes a conductive melt. Pass a direct current through it. Oxygen comes off at the anode, molten metal collects at the cathode, and there is no reagent and no consumable. The inputs are soil and electricity, and the Moon supplies one of them.

Fourteen kilograms of screened soil go in each day and four things come out: 1.80 kilograms of oxygen, 2.40 of iron silicon alloy, 9.20 of glassy slag for pavers and 0.60 of fines returned to the face. They add to 14.00, the first check to run on anybody's process diagram. Both rivals give something away. Hydrogen reduction of ilmenite runs at a gentler 900 degrees but attacks only the iron titanium oxide fraction, and one 2023 plant study found it pays only above 7.5 per cent ilmenite, more than this rim holds. Carbothermal reduction is further along, having run in vacuum at Johnson in 2023, but it needs a carbon loop and hands back no metal.

Melting carries the architecture's largest caveat. It has run on simulant and never on the Moon, and the anode has to survive oxygen at 1,600 degrees, which no material yet does for years. Anode life is the least certain number in the design, and it is why there are three cells.

WHERE PEOPLE GO WRONG

Ask anybody what a lunar base mines and the answer is polar ice. The Lunar Crater Observation and Sensing Satellite measured 5.6 per cent water by mass, plus or minus 2.9, in the plume it blew out of Cabeus, where the shadowed floors sit below the 110 kelvin that keeps ice for ever. Carry that error bar through: at 95 per cent confidence the grade runs from 11 per cent down to zero, so the most expensive water measurement ever made cannot rule out no ice. The first attempt to drill, PRIME 1 on the IM 2 lander, came to rest on its side in March 2025 and had ten hours instead of ten days. Oxygen in regolith is everywhere, measured eleven times over, and availability beats grade.

10.3 The number that decides whether any of it works

Every scheme for making oxygen on the Moon sounds good until somebody divides. A plant is two figures, the power it draws and the mass it makes, and the division that puts them together is unforgiving, because every kilowatt drawn on the Moon was carried there as reactor or as array. The design's own account, for each kilogram of oxygen, runs 8.4 kilowatt hours to melt the charge, 4.9 for electrolysis at the thermodynamic floor, 3.1 for overpotential, 0.7 for excavation, 1.4 to cool and compress the gas and 2.2 for losses and standby: 20.7 in all.

THE ONLY FORMULA IN THIS CHAPTER

$$e = (P \times t) \div m$$

e is the cost of the product in kilowatt hours per kilogram of oxygen. P is the electrical power the plant draws in kilowatts, t the time the output is counted over in hours, and m the oxygen made in that time in kilograms. There is no chemistry in this equation. It is a division, and it works for any plant that eats power and makes something. Both of its figures sit on the front page of every proposal ever written.

WORKED EXAMPLE 10.1 · WHAT A KILOGRAM OF OXYGEN COSTS IN POWER

The plant draws 14.0 kilowatts without stopping and makes 1.8 kilograms of oxygen a day. Schreiner's model rates a 400 kilogram, 14 kilowatt reactor at 1,000 kilograms of oxygen a year.

1. Find the oxygen the plant makes in a year.
   1.8 kg/day × 365 days = 657 kg
2. Find the energy it draws in the same year.
   14 kW × 8 760 h = 122 640 kWh
3. Divide, for the cost of a kilogram in this plant.
   122 640 kWh ÷ 657 kg = 187 kWh/kg
4. Now the model the plant was sized from, on the same 14 kilowatts.
   122 640 kWh ÷ 1 000 kg = 123 kWh/kg

ANSWER 187 kilowatt hours for every kilogram, and the figure is sound. It sits half again above the 123 of the model the plant is sized from, because it draws the same 14 kilowatts for two thirds as much oxygen, and it agrees with the 108 to 136 measured at Kennedy in vacuum. What does not describe this plant is 20.7. That is the same chemistry at ten to a hundred tonnes a year, where one cell's standing loss is divided among a hundred times the output. The design's losses line of 2.2 says loss is a ninth of the energy; here it is nine tenths.

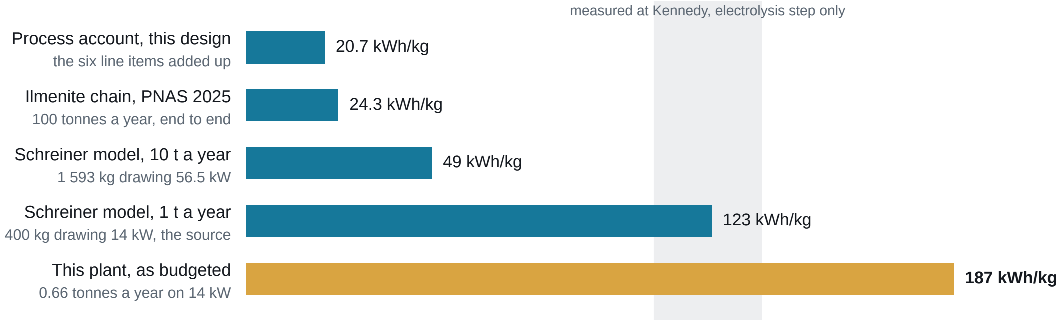

Figure 10.1. Energy per kilogram of oxygen, cheapest first. The list sorts by plant size, and the chemistry is identical all the way down it. What changes is how much oxygen the same standing heat loss is divided among. The only large scale run ever made, at Kennedy in late 2024, measured 108 to 136 for the electrolysis step alone.

So 187 is not a fault to be argued away, and the cure for it is not a cleverer cell but a bigger one. Wall area grows as the square of a reactor's size and the melt it holds as the cube, so four times the mass and four times the power buy ten times the oxygen, and the cost falls from 123 kilowatt hours a kilogram to 49.

The case for this plant is therefore freight and not energy. Fourteen kilowatts held for a year buys 657 kilograms of oxygen and 876 kilograms of metal nobody had to launch, and the margin is thin: 1.8 kilograms a day against a 1.65 kilogram gap is nine per cent. A figure in kilowatt hours per kilogram tells you nothing until you know how big the plant was.

QUESTIONS

1 Schreiner's model gives a 400 kilogram, 14 kilowatt reactor making 1,000 kilograms of oxygen a year and a 1,593 kilogram, 56.5 kilowatt one making 10,000. Find the cost of a kilogram in each, then say why this plant costs 187 on the same 14 kilowatts. APPLY

2 Explain why a small cell costs far more energy for each kilogram of oxygen than a large one, when the chemistry inside the two is identical. EXPLAIN

3 A drill confirms 5.6 per cent water by mass at a shadowed site nearby. Taking water as 88.8 per cent oxygen, how much must be dug for 1.8 kilograms of oxygen a day, and what would really decide between the two routes? EXTEND

4 Compute the oxygen fraction of the mare column in Table 10.1 from its rows, and show that it comes to about 42.9 per cent. CALCULATE

5 The plant makes 1.8 kilograms of oxygen and 2.4 kilograms of metal a day from 14 kilograms of soil. Explain where the rest of the mass goes, and why the four outputs must add to the input. EXPLAIN

6 The design's own account of 20.7 kilowatt hours per kilogram allots 2.2 to losses and standby. Using the 187 this plant actually costs, estimate what the losses line should read, and say what that implies about the anode and the insulation. EXTEND

CHAPTER ELEVEN

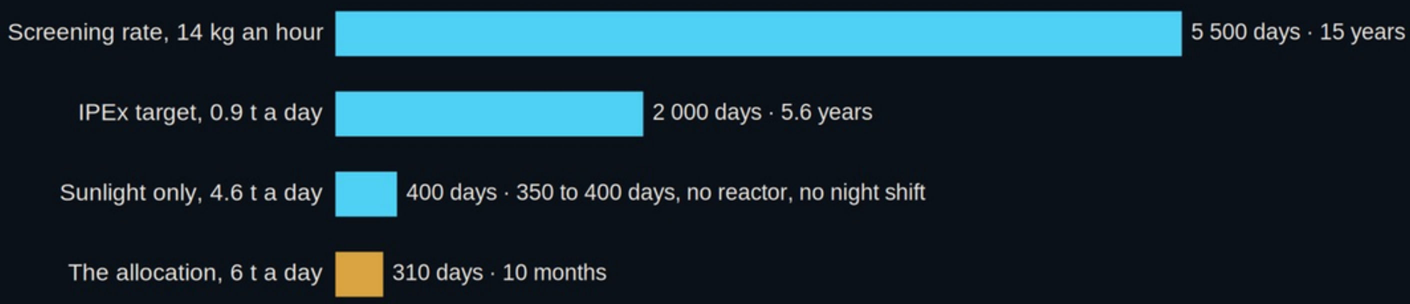

Construction

A lander touching down on bare regolith is the most destructive thing that can happen to a lunar base, so the first structure built is the one that lets the second spacecraft arrive. Everything after it is the same argument in other shapes: a laser, the ground, and a rate that has to be checked against the size of the job.

WHAT THIS CHAPTER WILL TEACH YOU

1. Explain why the landing pad is built first, and what a lander does to unprepared ground.
2. Describe how a laser turns loose regolith into a glass ceramic, and what forcing the cooling does.
3. Check a construction rate against the size of the job and against the energy available.
4. Say why a pad is inside what this material can do and a pressure shell is not.

155 m
APOLLO 12 TO SURVEYOR 3

1 360 °C
FULL MELT

5.6 m²
SLAB LAID A DAY

0 kg
BINDER SHIPPED

11.1 The pad is not a convenience

Every other system in this book defends the crew against the Moon. This one uses it.

A lander coming down on unprepared regolith is the most destructive event that can happen to a base: the exhaust plume does not push the surface aside, it digs into it and throws it. Plume models put the finest grains above two kilometres a second, against the 2.4 kilometres a second it takes to leave the Moon altogether. The harm is done by the grains that leave almost horizontally, because nothing slows them and nothing brings them down inside the site.

There is one measurement of the damage. Apollo 12 landed about 155 metres from the Surveyor 3 probe and sandblasted it, stripping paint and pitting the camera housing. The crew cut pieces off and carried them home, and the sizes of those pits bound the arrival speed between about 300 metres a second and two kilometres a second. That was a lunar module. The vehicle that delivers the next pressure vessel is far larger, and nobody has measured what one of those does.

A pad does not make the plume gentler. It takes away the loose material the plume works with and gives the lander a surface that will not collapse under a leg. So the order of work is forced rather than chosen: survey, grade flat, roll until the voids close, melt a slab, and only then bring anything else down beside the base, 65 metres away.

11.2 Melting the ground into ceramic

Regolith is mostly silicates and metal oxides, and it does not burn. Heat it and the grains begin to weld to one another near 1,180 degrees Celsius; past about 1,360 it melts outright and freezes on cooling into a glass ceramic that is strong in compression. Both figures come from laser tests on a south pole simulant, which also found melting about ten times more energy efficient for each unit consolidated than sintering. So the head melts rather than sinters, and nothing is added to the ground: no cement, no water, no binder, nothing shipped for it.

The sequence for one course is grade, roll, melt, wait, inspect. The rolling matters as much as the beam, because melting an uncompacted surface leaves a crust bridging cavities, and the crust breaks the first time a lander puts weight on it. The beam takes the top 30 millimetres past 1,360 degrees, one course at a time, and then the course has to cool, which in vacuum means radiating to the sky and conducting downward into cold soil. Both are slow, and slow is what makes the slab strong.

WHERE PEOPLE GO WRONG

Sintered regolith gets called lunar concrete, and the name does real damage. Concrete is aggregate held together by a cement paste that needs water to cure. This is glass: melted, then frozen, with nothing added. In compression it holds 20 to 45 megapascals against 20 to 40 for concrete, near enough the same. In tension it holds 2 to 4 megapascals, about a tenth as much. A pad, a road, a berm and a footing are loaded in compression only, so everything through increment three sits well inside what the material can do. A pressure vessel is the other case entirely, because the air inside pulls the wall apart. That is why the printed shell is a trial with iron in it, and not a plan.

WORKED EXAMPLE 11.1 · CAN A 9 KILOWATT BEAM LAY 11 SQUARE METRES A DAY?

A course 30 millimetres deep in regolith of about 1,500 kilograms per cubic metre, so a square metre is 45 kilograms, and 0.86 kilowatt hours to take a kilogram from lunar cold to a melt. The beam is 9 kilowatts and the design claims 11 square metres a day.

1. Find the energy one square metre of course needs.
   45 kg × 0.86 kWh/kg = 38.7 kWh, call it 39
2. Find what the beam delivers in a whole day without a break.
   9 kW × 24 h = 216 kWh
3. Divide, to see how much slab that pays for.
   216 kWh ÷ 38.7 kWh = 5.6 m^2
4. Then ask what the claimed rate would take.
   11 × 38.7 = 426 kWh, and 426 kWh ÷ 24 h = 17.8 kW

ANSWER About 5.6 square metres a day, half the rate the design claims, and no engineering closes that gap, because every joule in step one is thermodynamics: the heat capacity and the latent heat of the material. Keeping the claimed rate would mean holding 17.8 kilowatts on the beam, twice the average draw the design allows the whole construction system. So this book carries 5.6 square metres a day, and the apron takes 68 days rather than 35. Even 5.6 is a ceiling, since it assumes every joule from the socket ends up in the melt, when a laser of this class turns 10 to 20 per cent of its input into beam.

11.3 A rate is not a schedule

The build order lists eight jobs and gives each a duration, and every duration is a job size divided by a rate. A student who can do that division can audit the whole programme in an afternoon. The trap is that a rate belongs to a task and not to a machine. The same excavator appears twice in this design, at 14 kilograms an hour and at 6 tonnes a day, and those two numbers describe different work. Put the largest job through the first of them: 1,850 tonnes of cover at 14 kilograms an hour is 132,000 machine hours, which is 5,500 days, or fifteen years.

Fourteen kilograms an hour is a feed rate, and a gentle one: the oxygen plant eats 14 kilograms a day, so one seven hour shift a week keeps its hopper ahead. Machines of this class dig far faster. NASA's ISRU Pilot Excavator is designed for 42 kilograms an hour, and the RASSOR 2.0 breadboard moved 2,720 kilograms an hour while its drums turned. The build order's 6 tonnes a day is 250 kilograms an hour, so the hard part is not digging but holding it around the clock: that same excavator moves 10 tonnes in 11 days all in, about 0.9 tonnes a day.

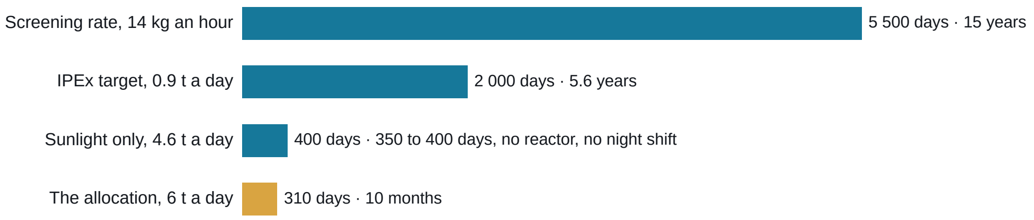

Figure 11.1. How long the 1,850 tonne cover takes at four rates. Only the bottom bar is in the build order, and the third is what waiting for the Sun would cost on this rim, lit 77 to 88 per cent of the year.

Table 11.1. What gets built, in the order it has to be built, and the hazard each job answers.

Structure	Increment	Size and time	What it answers
Graded pad bed	1	380 m², 8 days	Nothing can land safely until it exists
Sintered apron, 22 m across	1	380 m², 68 days	The slab itself, laid at 5.6 m² a day
Blast berm, 29 m outer	2	1 400 pavers, 36 days	The low angle fan the apron cannot catch
Module cover, 2 m of regolith	**2**	**1 850 t, 310 days**	**Cosmic rays. The largest job, and not begun**
Graded roads, four routes	3	2 100 m², 247 days	Dust ground out by wheels, and rover wear
Printed shell trial	4	18 m², 60 days	Whether habitat volume is launched or made

The apron and the roads are recomputed here, at the 5.6 square metres a day the beam can pay for and at the design's own road rate of 8.5, against the 35 and 190 days it prints. The equipment pads, 260 square metres, take 46 days rather than 24. The skirts at the foot of each hull are not budgeted at all.

The last line of the table is what the programme is for. Of the 48.3 tonnes this base lands, 28.7 are pressure vessels, and that cannot be argued down without taking volume from the crew. The trial article, 18 square metres of printed shell with iron drawn from the plant's own alloy, is the experiment that decides whether the next base lands its rooms or makes them.

QUESTIONS

1. Use T = M ÷ (R × H) on the 1,850 tonne cover, at 14 kilograms an hour and then at 6 tonnes a day. What rate do the 310 days demand? APPLY

2. Explain why a melted course is left to cool at its own pace, and what happens to the slab if the cooling is forced. EXPLAIN

3 The same material makes a pad that survives a landing and would not make a pressure shell. Explain why, and say what the increment four trial has to show. EXTEND

4 Rework worked example 11.1 for a course 50 millimetres deep, and say what happens to the schedule for the apron. CALCULATE

5 The four graded routes total 2,100 square metres at 8.5 square metres a day. Show that the design's 190 days corresponds to the 11 square metre slab rate rather than the road rate, and give the corrected figure. CALCULATE

6 A laser of this class turns 10 to 20 per cent of its electrical input into beam. Say what that does to the 5.6 square metres a day, and what happens to the rest of the energy. EXTEND

the air steps are the ones people overestimate: one oxygen molecule out for each carbon dioxide fixed

01 → 02 → 03 → 04 → 05 → 06

01 Crew air
4.2 kg of carbon dioxide a day

02 Greenhouse bay
22 m² on six tiers, taking 41 g a day

03 Oxygen back
0.03 kg a day, one per cent of demand

04 Transpiration
2 litres per square metre a day

05 Condensers
2.1 kg a day to the cabin water loop

06 Harvest
4.1 kg fresh a week, 30 g dry a day

the bay is properly part of the water loop and barely part of the air loop

Food

Six tiers of trays in a five metre module grow about one per cent of what the crew eat. That is the design working rather than the design failing, and the argument for it is a cost per calorie rather than a slogan about closing the loop.

WHAT THIS CHAPTER WILL TEACH YOU

1. Say why the staples are shipped and the salad is grown, and what settles the split.
2. Explain why one room of 54 cubic metres draws as much power as the oxygen plant.
3. Size the growing area a share of one person's diet needs, from a yield and an energy density.
4. State what the bay gives back to the air and water loops, and what it gives the crew that is not food.

22 m²
GROWING AREA

4.1 kg a week
HARVEST

14 kW
LIGHTING LOAD

about 1 per cent
CALORIES GROWN

12.1 The calculation everybody skips

The case for growing food in space is usually made on closure: a loop that closes needs no resupply. The number that decides it is the cost of a delivered calorie, and for staples it points the other way.

A kilogram of wheat grain holds about 3,300 kilocalories that keep for years in a sealed bag, and shipping it costs one kilogram of landed mass. Growing the same kilogram costs light, water, volume and crew hours, for months. In NASA's Biomass Production Chamber, wheat under 510 to 930 micromoles of light per square metre per second made 23 to 40 grams of dry biomass per square metre a day and took 77 to 86 days to mature, and only part of it is grain. NASA's summary of the programme is that at 40 moles of light a day, roughly 50 square metres of crop feeds one person. This bay is 22 square metres and four people live off it, so the design does not try to feed them.

The design calls the harvest one per cent of what they eat, worth checking both ways. By energy, 120 kilocalories a day against four people near 12,000 is one per cent. By dry mass, 4.1 kilograms of fresh salad a week is only about 30 grams of dry matter a day, against the 2.72 kilograms of dry food the crew land each day, which is 1.1 per cent. The two agree.

Table 12.1. Ship it or grow it, crop by crop. Yields are dry biomass in NASA's chamber, and the energy values are for the raw crop.

Crop	Dry yield, g/m^2 a day	kcal per kg, raw	How long it keeps	Decision
Wheat and rice, staple	23 to 40	3 300 to 3 600	Years, sealed	Ship it
Soy and pulses, protein	10 to 16	3 500 to 4 500	Years	Ship it
Potato, staple	22 to 33	770	Months	Ship it
Lettuce, chard, kale	**6 to 8**	**150 to 170**	**Days**	**Grow it**
Tomato and pepper	**13 to 20**	**180 to 310**	**Weeks**	**Grow it**
Radish, microgreens, herbs	**not tested**	**160**	**Days**	**Grow it**

What settles the last column is the fourth, not the second. Wheat and pulses keep for years sealed, so a kilogram shipped is a kilogram eaten however long it waits, while lettuce keeps for days and its vitamins fall away the whole time. The crops worth growing are the ones with hardly any calories in them.

0.22 kg a week Herbs, out of all proportion. The week's 4.1 kilograms is 1.55 of leafy greens, 1.05 of tomato and pepper, 0.68 of radish and microgreens, 0.60 of dwarf bean and 0.22 of herbs. The smallest line is the one the crew would miss first: it is the only one that can change what a ration tastes like.

12.2 Manufacturing sunlight

Of everything on this base, only the plants need sunlight made for them. There is no daylight inside a pressure vessel under two metres of soil, so every photon a leaf receives was electricity a moment earlier, and the bay draws 14 kilowatts to make them.

Chlorophyll absorbs most strongly in the deep red near 660 nanometres and the blue near 450, and least in the green, which is why leaves look green. That is the usual explanation for red and blue lamps, and it is the wrong one: McCree's measurements across 22 crop species put the quantum yield of an absorbed green photon close to that of a red one. The case is electrical. A 660 nanometre diode delivers up to 5.5 micromoles of photons for each joule, the best of any colour, where a green one manages 1.9 and a white one loses part of its energy in the phosphor. Red and blue buy the most photons per watt, and their sum is magenta.

How much light, and for how long, is not settled. The design runs an 18 hour photoperiod; Veggie runs 16 hours at about 200 micromoles per square metre per second and the Advanced Plant Habitat goes to 1,000. The 14 kilowatts is an allocation built on a figure nobody has fixed. A little green is added anyway, one diode in sixteen on Veggie, because under red and blue alone nobody can tell a sick plant from a thirsty one.

WORKED EXAMPLE 12.1 · WHAT A KILOGRAM OF SALAD COSTS IN ELECTRICITY

Lighting of 14 kilowatts on an 18 hour photoperiod, and a harvest of 4.1 kilograms of fresh mass a week carrying 120 kilocalories a day. Take 1 kilocalorie as 4.184 kilojoules and 1 kilowatt hour as 3,600 kilojoules.

1. Energy into the lamps in a day, then in the harvest's week.
   14 kW × 18 h = 252 kWh a day, and 252 kWh × 7 d = 1 764 kWh
2. Divide by the harvest for the electrical cost of a kilogram of crop.
   1 764 kWh ÷ 4.1 kg = 430 kWh/kg
3. Put the food energy in that harvest into the same units.
   120 kcal × 7 d = 840 kcal, and 840 × 4.184 kJ = 3 515 kJ, which is 0.98 kWh
4. Compare the electricity in with the food energy out.
   1 764 kWh ÷ 0.98 kWh = about 1 800

ANSWER About 430 kilowatt hours of light for each kilogram of salad, and about 1,800 joules of electricity for every joule of food energy that comes back. Only the lamps are counted: pumps, fans and the heat all 14 kilowatts becomes are on top. The ratio is this ugly because the crops are mostly water and were chosen for vitamins and flavour rather than calories. Wheat at 3,300 kilocalories a kilogram would score far better and still lose to the freighter, because it wants two to five times the daily light of lettuce and 77 to 86 days to get there.

SIZING A GROWING AREA

$$A = (f \times E) \div (y \times c)$$

A is the growing area needed in square metres, f the share of one person's food energy to be grown, E one person's food energy in kilocalories a day, y the fresh yield in kilograms per square metre a day, and c the energy density as eaten in kilocalories per kilogram. For this bay, 4.1 kilograms a week over 22 square metres gives y = 0.027, and the mix is about 205 kilocalories a kilogram, so 0.027 × 205 = 5.5 kilocalories per square metre a day. One person at 3,000 kilocalories would need 545 square metres. NASA's 50 square metres a person is the same equation with wheat at 3,300 kilocalories a kilogram in place of salad at 205.

12.3 Into the loops, and into the crew

The bay is properly part of the water loop and barely part of the air loop. Plants transpire nearly all the water they take up, about 2 litres per square metre a day for lettuce, so tens of litres pass through the condensers daily, most of it straight back to the nutrient solution, with 2.1 kilograms a day credited to the cabin loop.

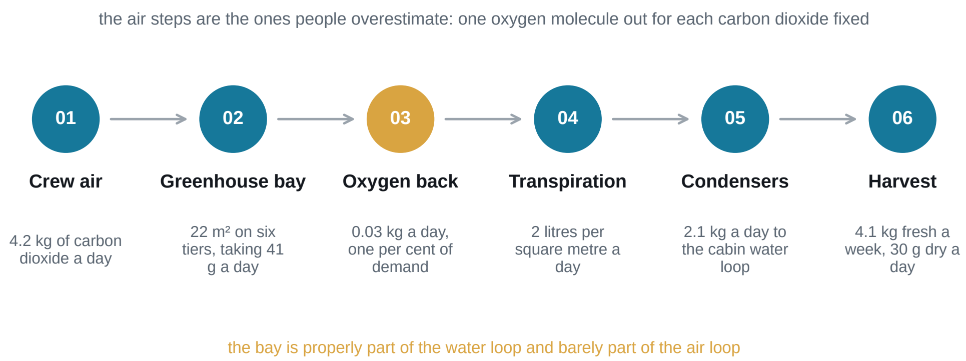

Figure 12.1. What the bay exchanges with the rest of the base. The air steps are the ones people overestimate: photosynthesis releases one molecule of oxygen for each molecule of carbon dioxide it fixes, so 0.03 kilograms of oxygen a day, which is 0.94 moles, means 41 grams of carbon dioxide taken up.

WHERE PEOPLE GO WRONG

Popular accounts hand the greenhouse a quarter of the base's air. Work it instead. The harvest is 4.1 kilograms of fresh mass a week, mostly water, so about 30 grams of dry matter a day, and photosynthesis releases about a gram of oxygen for each gram of dry matter. The bay therefore makes 0.03 kilograms of oxygen a day against the 3.36 four crew breathe, and takes up 41 grams of carbon dioxide against the 4.2 kilograms they make. One per cent. The error is to picture a greenhouse as a volume of green leaf, when what sets the oxygen is the rate at which dry matter is built, and salad builds very little. NASA's figure is 20 to 25 square metres of crop for one person's oxygen, so this crew would need 80 to 100. The oxygen here comes out of the regolith.

What the bay gives the crew that is not food is harder to weigh and easier to lose. Fresh taste, colour, a task with a visible result, and something alive in a machine. Crews on long missions report that growing things matters out of all proportion to the calories, and the herbs line in the harvest, the smallest by mass, is the one that shapes what the rations taste like. That is not sentiment; food acceptance falls on a monotonous diet, and a crew that stops eating well stops working well.

QUESTIONS

1 The harvest is 4.1 kilograms fresh a week, about 30 grams of dry matter a day, carrying 120 kilocalories. The crew land 2.72 kilograms of dry food a day and eat about 12,000 kilocalories. Work out the bay's share both ways. APPLY

2 An absorbed green photon drives photosynthesis nearly as well as a red one. Explain why the lamps are red and blue even so, and why a few green diodes are added. EXPLAIN

3 The chamber grew 6 to 8 grams of dry lettuce per square metre a day. What would 22 square metres give, and what does the gap with the 30 grams budgeted tell you? EXTEND

4 Use $A = (f \times E) \div (y \times c)$ to find the area needed to grow a tenth of one person's 3,000 kilocalories in salad at 205 kilocalories a kilogram and 0.027 kilograms per square metre a day. CALCULATE

5 Explain why the greenhouse is the largest load the base can shed, and what happens to the crops if it is shed for the 66 hour dark spell of chapter 5. EXPLAIN

6 The bay credits 2.1 kilograms of water a day to the cabin loop. Where does that water come from, and what does the crew have to put back into the nutrient solution to keep the balance? EXTEND

Medicine

A crew of four, two of them trained as medical officers, and a hospital four days away. Most of what this chapter plans for will not happen to any one person. Some of it will happen to somebody.

WHAT THIS CHAPTER WILL TEACH YOU

1. Name the numbers in this chapter that were measured at one sixth gravity, and say why the answer is none of them.
2. Explain why a ride home of four days at best turns stabilise and evacuate into treat in place.
3. Turn an incidence rate into an expected number for a different crew and a different duration, and say what that arithmetic quietly assumes.
4. Judge what 428 kilograms of medical mass buys, and name what it cannot treat at all.

74:59	4 days	180 days	428 kg
HUMAN RECORD AT ONE SIXTH G, HOURS AND MINUTES	RIDE HOME, BEST CASE	ROTATION LENGTH	MEDICAL MASS LANDED

13.1 Everything here was measured somewhere else

On 14 December 1972 two men closed the hatch of a lunar module and lifted off. Eugene Cernan and Harrison Schmitt had been on the surface for 74 hours and 59 minutes. That is the entire human experience of sustained one sixth gravity: two people, three days, and nothing since.

So read the chapter knowing what it is made of. Every bone figure in it, every muscle figure and every eye figure, was measured on somebody weightless in orbit and then carried across. A 2017 review of human work in partial gravity found 43 studies and 285 participants, and reported that none of them investigated long term effects: parabolic flights, tilt rigs and suspension, minutes to a fortnight. Ask which number here was measured at 0.166 g. None of them were.

The second hard fact is the ride home. Apollo 11 lifted off from the surface on 21 July 1969 and splashed down on 24 July, about three days, waiting for no launch window. NASA's surface mission profile allows a crew up to five days in lunar orbit before Orion turns for Earth. This design plans on four days at best, seven to ten once a window has to be waited for. Nobody has ever been evacuated from the Moon for a medical reason, so both figures are estimates.

That reverses the medical question. On the station the job is to hold a patient steady until the capsule undocks. Here, with definitive care four days away at best, the job is to treat somebody in place long enough that evacuation becomes optional. Two of the four crew are trained as medical officers, about 40 hours of training, as on the station. Neither is a surgeon, and the base has no surgical capability.

> **74:59** Apollo 17's surface stay, and the whole human record at this gravity. A rotation asks for 180 days, or 4,320 hours, so the design extrapolates by a factor of about 58. Increment three begins to answer it: continuous densitometry through a full rotation, the first partial gravity data anyone will have.

13.2 A small number is not a zero

The risk list is not the one intuition offers. Trauma is rare. What fills a rotation is a set of changes that are slow, predictable and cumulative, along with complaints that are trivial at home and are not trivial here. A dental abscess is an afternoon on Earth. Four days from a dentist it can end a rotation.

The bars are an allocation rather than a forecast, built the way NASA's Integrated Medical Model works, from station incidence scaled to four people and 180 days. Two caveats sit under every bar: the rates come from crew a few hours from a hospital, and none of them were measured at a sixth of a gravity. Now read the bottom bar. Serious trauma or an acute abdomen is allocated 0.05 events in a rotation, which reads as nothing at all. Work out what it means across a whole programme before you read on.

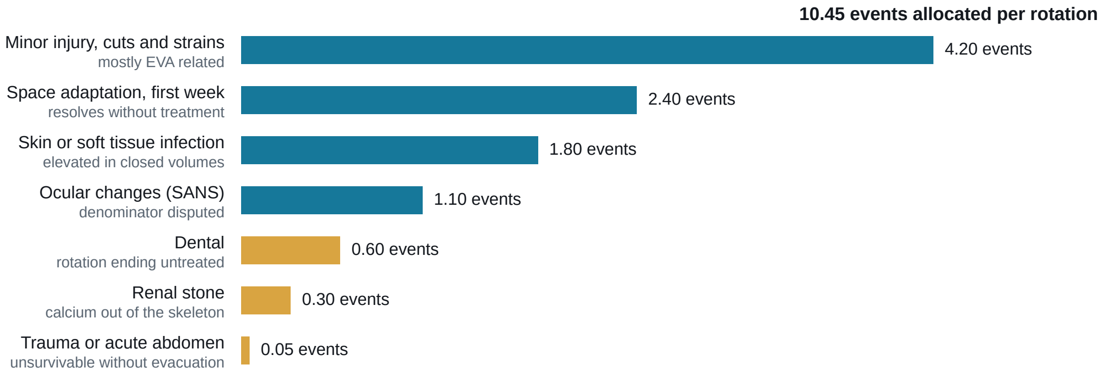

Figure 13.1. Expected medical events in one 180 day rotation across a crew of four, as the design allocates them. The bars add to 10.45. The station's own clinical record, 46 crew over 20.6 flight years at 3.4 events a flight year, scales to 6.7 events for this crew and this rotation, so the allocation is about half as much again and the two are plainly not counting the same list. Read the amber bars rather than the top one: everything above them is uncomfortable, and everything below decides whether a rotation ends early.

THE ONLY FORMULA IN THIS CHAPTER

$$E = r \times n \times t$$

E is the expected number of events over the whole exposure, r the incidence rate in events per person per year, n the number of people exposed, four on this base, and t the exposure in years, so a 180 day rotation is 0.493. E is an average and not a count. It can be smaller than one, and a number smaller than one is not a promise that nothing happens: it is the figure you multiply by the number of rotations. If events arrive independently at a steady rate, the chance that an exposure passes with none of them at all is $1 \div e^{E}$, and the chance of at least one is whatever is left of 1.

WORKED EXAMPLE 13.1 · THE EVENT THAT WILL PROBABLY NOT HAPPEN

The shortest bar on the chart, serious trauma or an acute abdomen, at 0.05 expected events in one 180 day rotation across four crew. The design does not say how many rotations there will be. Take ten, and change the number if you prefer: the method does not alter.

1. Turn the rotation figure back into a rate per person per year, so it can be carried to any crew and any duration.
   t = 180 d ÷ 365 d/y = 0.493 y, and r = 0.05 ÷ 1.97 person years = 0.0253 per person per year
2. Run the rate forward over the whole programme. Ten rotations is 1,800 days, which is 4.93 years, and the crew is still four, so the exposure is 19.7 person years.
   E = 0.0253 × 19.7 person years = 0.50 events
3. Turn the expectation into a probability. Take the events as independent and the rate as steady, and the chance of the programme passing with none of them at all is one divided by e raised to E.
   $1 \div e^{0.50} = 1 \div 1.649 = 0.607$
4. None and at least one are the only two outcomes there are, so their chances add to 1.
   1.000 take away 0.607 = 0.393, about two chances in five

ANSWER About two chances in five that a serious trauma or an acute abdomen happens somewhere in the programme, from a line that reads 0.05 and looks like nothing. That is the argument for the 28 kilograms of trauma stabilisation, which will most likely be landed, stowed, inspected for five years and never opened. It is also the argument against reading too much into the answer. Those 28 kilograms buy time and nothing else, four days of it at best, because the base has no surgical capability. And the 0.05 came from a station crew a few hours from a hospital, which is not the crew this arithmetic has just been applied to.

WHERE PEOPLE GO WRONG

The chart gets read as a forecast. It is not. An allocation of 4.2 minor injuries does not mean four injuries will happen: it is a rate multiplied by four people and by 180 days, and the count that actually arrives could be one or eight. At the other end of the chart the mistake is worse, because 0.05 gets read as zero. And where the evidence for a countermeasure is called strong, it is strong for orbit. Nothing in this chapter has been tested at a sixth of a gravity.

13.3 What 428 kilograms buys

Bone is the certainty on that chart. In weightlessness crews lost bone mineral density at about one to one and a half per cent a month on the earlier station devices; with the advanced resistive exercise device and good nutrition the loss at hip and spine falls to 0.3 to 0.5 per cent a month. It does not stop. At 0.166 g the skeleton is genuinely loaded, and whether that arrests the loss, slows it or barely touches it has never been measured on anyone.

That is what 310 of the 428 kilograms are for. An hour and a half of resistive work and half an hour on the ergometer, inside a two and a half hour slot once rigging and stowage are counted: across four people, eight to ten crew hours a day against the twenty the base allots to work. Everything that treats comes to 118 kilograms.

Table 13.1. The medical kit, by category and allocated mass.

Category	Contents	Mass	Note
Diagnostics and monitoring	Ultrasound, portable analyser, vitals suite, dosimeters	50 kg	The ultrasound is the workhorse: no consumables, no radiation
Trauma stabilisation	Airway, haemorrhage control, splinting, sutures	28 kg	Stabilise for a return, not definitive repair
Pharmacy, 180 day supply	Analgesia, antibiotics, cardiac, psychiatric	34 kg	Shelf life under radiation is the limit, not quantity
Dental	Temporary filling, extraction, analgesia	6 kg	An afternoon on Earth, a rotation ending here
Exercise hardware	Resistive flywheel, cycle ergometer	310 kg	The heaviest medical item on the base, and the most used
Total		**428 kg**	**Three quarters of it prevents rather than treats**

The masses are allocations. NASA's Exploration Medical Capability element has not published a kit mass for a surface stay of this length, so there is no flown figure to check them against. Monitoring, 8 kilograms of dosimeters and sleep and activity sensors, sits inside the first row.

QUESTIONS

1. A 2024 systematic review reported optic disc swelling in about three quarters of long duration crew imaged after flight. NASA's own evidence report puts the clinical incidence of the syndrome across the astronaut corps at 16 per cent. Turn both into an expected number for a crew of four, and set them against the 1.1 ocular events the chart allocates. APPLY

2. Explain why a 2.56 second round trip to Earth is acceptable for a consultation and unusable during a procedure, and describe the model the base adopts instead. EXPLAIN

3. Of the 428 kilogram kit, 310 kilograms prevent and 118 kilograms treat. Argue whether that is the right split for a base four days from definitive care, and say what would have to change to move mass the other way. EXTEND

4. Rework worked example 13.1 for a programme of twenty rotations, and say what the answer implies for whether the base should carry surgical capability. CALCULATE

5. Bone loss at 0.3 to 0.5 per cent a month does not stop in orbit. Estimate the loss over a 180 day rotation at each rate, and say why neither figure can be trusted here. CALCULATE

6. The pharmacy is limited by shelf life under radiation rather than by quantity. Say where on the base the drugs should be stored, and what that decision costs in convenience. EXTEND

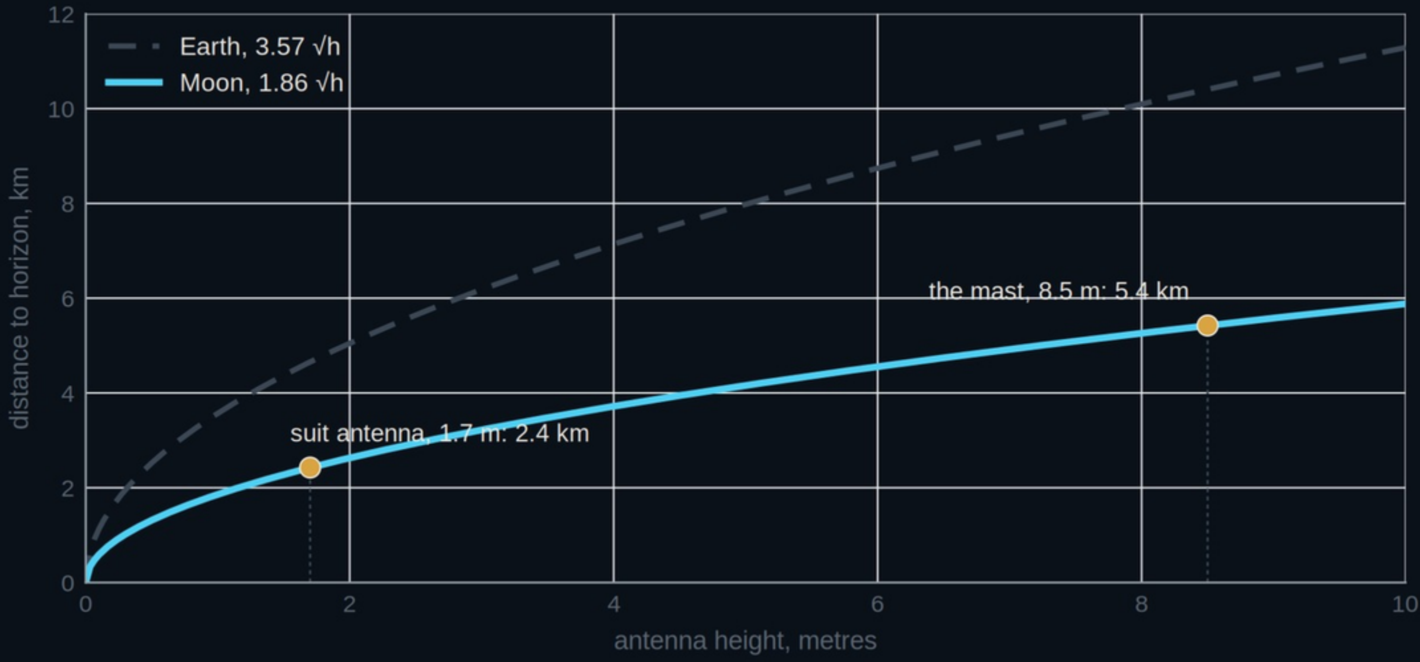

Communications

Two constraints run this subsystem and both of them are geometry. Light takes 1.28 seconds to reach the Moon, and anything below the horizon of an airless body is not faint, it is gone. No amount of bandwidth touches either one.

WHAT THIS CHAPTER WILL TEACH YOU

1. Work the distance to the horizon from a body's radius and an antenna height, and say why the constant is 1.86 on the Moon against 3.57 on Earth.
2. Explain what a 2.56 second round trip does to anybody trying to drive a machine on the surface from the ground.
3. Say why a relay in a low lunar orbit is in sight for minutes rather than hours, and what the network does with the rest of the time.
4. Separate the parts of this network that are geometry from the parts that are engineering, and say which of the two a bigger amplifier can fix.

2.56 s
ROUND TRIP TO EARTH

57 per cent
EARTH ABOVE THE RIM

5.4 km
HORIZON FROM THE MAST

340 kg
NETWORK LANDED

14.1 Nothing here is a bandwidth problem

Drive a rover on the Moon from a desk on Earth and you are driving blind. The Moon is 384,400 kilometres away on average, light takes 1.28 seconds to cross that, and a command and the picture that shows what it did take 2.56 seconds between them.

At a construction rover's half a metre a second, the machine covers 1.3 metres of unseen ground on every command. That is about the size of the boulder that ends the mission. No amount of bandwidth touches that number and no engineering decision changes it. It is why autonomy on this base is a communications requirement rather than a separate subsystem, and why the network is sized by geometry instead of by data rate.

The second constraint is the horizon, and there is no ionosphere to bounce a signal off. Anything below the horizon is not faint, it is gone. The surface cell has 51 decibels of margin at two kilometres, which is a great deal of spare power, and every decibel of it is worthless the moment the far end drops out of sight. Coverage is bought with height, not with amplifiers.

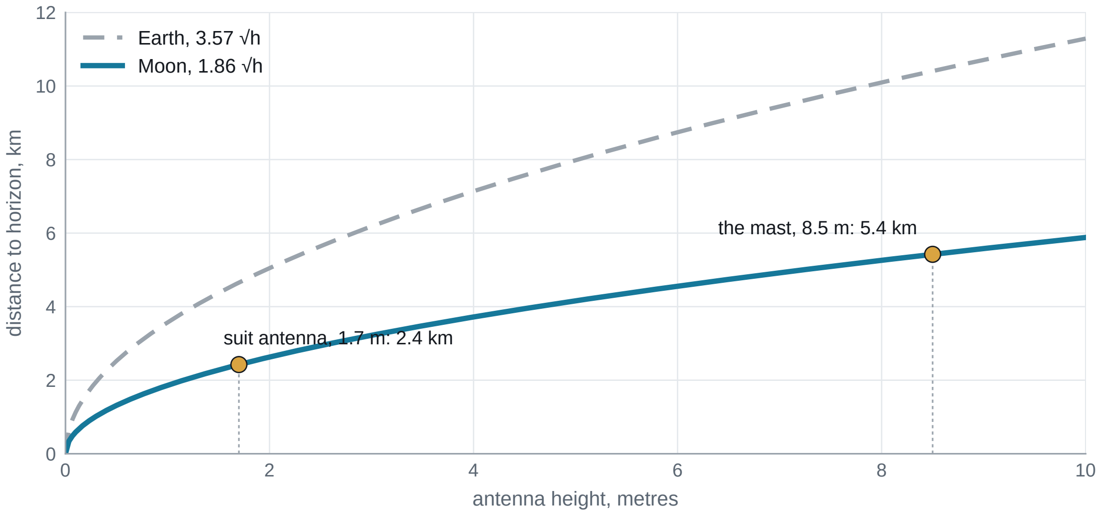

Figure 14.1. Distance to the horizon against antenna height, on a smooth sphere. The Moon's smaller radius costs about half the range for the same mast: 1.86 times the square root of the height against 3.57 on Earth. An 8.5 metre mast sees 5.4 kilometres and a suit antenna at 1.7 metres sees 2.4. Because it is a square root, range gets dearer as you buy it: the first kilometre costs 0.29 metres of height and the second costs 0.87 more, three times as much. At the Shackleton rim the crater wall and the ridge cut all of these figures shorter.

14.2 The square root that sets the mast

Everything after this follows from one line of school geometry. Draw the tangent from the top of the mast to the sphere, apply Pythagoras at the right angle where it touches, and throw away the term in the square of the height, because the height of a mast is nothing beside the radius of a moon. What is left is the horizon formula, and once the Moon's radius goes in it collapses to a single constant.

THE ONLY FORMULA IN THIS CHAPTER

$$d = \sqrt{2 R h}$$

d is the distance from the antenna to its horizon, R the radius of the body, about 1,740 kilometres for the Moon, and h the height of the antenna above the surface. The term in h squared is dropped because a mast is a millionth of a radius. Put R in and the formula becomes one constant: on the Moon, 1.86 times the square root of the height in metres, in kilometres. On Earth the same formula gives 3.57. The Moon is a smaller ball, so from the same mast you see about half as far, and the whole of the rest of this chapter follows from that.

WORKED EXAMPLE 14.1 · WHY THE MAST IS 8.5 METRES

The Moon's radius of 1,740 kilometres, a suit antenna 1.7 metres above the ground, and a rover that has to keep working 5 kilometres out from the habitat.

1. Work the constant out once, in the units the answer will be read in: heights in metres, distances in kilometres.
   $\sqrt{2 \times 1\,740\,000 \text{ m}} = 1\,865$ m per $\sqrt{\text{m}}$, which is 1.86 km per $\sqrt{\text{m}}$
2. Take the suit antenna first. A crew member standing on level ground carries it about 1.7 metres up.
   $d = 1.86 \times \sqrt{1.7} = 1.86 \times 1.30 = 2.4$ km
3. Now the mast. Height is the only term the design controls, and it sits under a square root.
   $d = 1.86 \times \sqrt{8.5} = 1.86 \times 2.92 = 5.4$ km
4. Check it against the job, and against the mast that was rejected.
   5.4 km covers the 5 km traverse; a 3 m mast gives $1.86 \times 1.73 = 3.2$ km and does not

ANSWER An 8.5 metre mast sees 5.4 kilometres on a smooth sphere and a 3 metre mast sees 3.2. The extra 5.5 metres of mast buys 2.2 kilometres of horizon, and that is what the 148 kilograms of mast and pointing gear is for. Two warnings come with the answer. The formula assumes a smooth sphere, and at the rim the crater wall and the ridge cut the real horizon shorter in most directions. And the number that matters on site today is not 5.4 kilometres at all: the excavation face is 84 metres from the habitat, which anything on the mast can see over. The mast is tall for the traverses that have not happened yet, not for the base that is there now.

WHERE PEOPLE GO WRONG

Coverage gets treated as a power problem. The surface link carries 51 decibels of margin at two kilometres, so the transmitter could be turned down by a factor of a hundred thousand and the link would still close. It does not fail from distance. It fails when the far end drops below the horizon or behind the 1.5 metre blast berm around the landing pad, and no amount of transmitter power reaches around either of those. The second half of the mistake is the smooth sphere: 5.4 kilometres assumes one, and the ridge is not one.

14.3 An ordinary radio, and a network built to wait

The radio hanging on that mast is deliberately ordinary. Everything on the surface joins one cellular network, LTE now and 5G later, because a bespoke radio would be lighter and would still be the wrong answer: cellular chipsets are made in the hundreds of millions, cellular networks solved the moving terminal in the 1980s, and 5G is a change of radio head on the same core. Nokia flew such a network to the surface on IM 2 in March 2025. The lander came to rest on its side, the system had about 25 minutes of power, it woke up and returned data, and no call was ever placed on the Moon. That is thin heritage, and it is more than any other tier of this network has.

Above the surface cell the network reaches Earth two ways. A Ka band dish looks at Earth directly when Earth is above the rim, which at this site is about 57 per cent of the time and never more than a few degrees up, so the dish needs a gimbal and a clear bearing along it. When Earth is down, a relay in a low lunar orbit takes the traffic, and a relay is in sight for minutes rather than hours: at most 12.4 minutes from a 100 kilometre circular orbit, and only from a pass that goes directly overhead. Every other pass is shorter, and the orbit takes 117.8 minutes to come round again, so the planning figure is about ten minutes of contact once every two hours.

So it is built never to wait for Earth. TCP, the protocol the internet runs on, reads a missing acknowledgement as congestion and backs off, which over a 2.56 second round trip and outages measured in hours is a protocol that spends its life waiting. In its place is the store and forward Bundle Protocol, where each hop takes custody of the data before the last one lets go. Everything is written to an eight terabyte store at the mast and forwarded when a path exists. A transfer broken by an outage resumes rather than restarting, and the crew see a mailbox rather than a connection.

All of it lands as 340 kilograms drawing 0.9 kilowatts, which is 3.3 per cent of the habitat load. The mast alone is 148 of those kilograms and the cabling is 34, which is the line in any mass budget most likely to be wrong.

> **12.4 minutes** Longest a relay in a 100 kilometre orbit stays in sight, and only from a pass directly overhead. Computed from Kepler's law with the Moon's gravitational parameter and radius; the design said a few tens of minutes.

Table 14.1. What each kind of traffic does when a path goes away.

Traffic	Priority	Path	Behaviour on an outage
Crew voice and video	1	Relay, else direct	Degrades to voice, then to text
Caution and warning	1	Both, duplicated	Never queued, always duplicated
Command and control	2	Relay preferred	Queued, executed on arrival
Science and imagery	4	Whichever is up	Queued at the mast, forwarded later
Housekeeping telemetry	4	Whichever is up	Summarised, then forwarded

Priority is a number the Bundle Protocol carries with every bundle it holds. Only the top two rows have to be alive. Everything else is a file that arrives when it arrives, which is the difference between a network built around a delay and one built around a connection.

QUESTIONS

1 A ridge top mast has to keep a rover working 12 kilometres away on the network. Using $d = 1.86\sqrt{h}$ with h in metres, find the mast height needed if the rover's own antenna is on the ground, and then if the rover carries a 1.7 metre antenna of its own. Say what the two answers show about buying range with height. APPLY

2 Explain why a 2.56 second round trip makes driving a machine from the ground unsafe while leaving a voice conversation perfectly workable, and say what the base does instead. EXPLAIN

3 The 57 per cent Earth visibility figure comes from one point on the Shackleton rim in a published study. Say what could move it and in which direction, and what the design would have to change if the true figure at the mast turned out to be 40 per cent. EXTEND

4 A relay in a 100 kilometre orbit has a period of 117.8 minutes. Show from the geometry of the horizon why the longest pass is 12.4 minutes, and say what a 200 kilometre orbit would change. CALCULATE

5 The surface cell has 51 decibels of margin at two kilometres. Explain what a decibel is, what a factor of a hundred thousand in power means for the link, and why none of it helps when a rover drives behind the berm. EXPLAIN

6 The Bundle Protocol gives every message a priority and a custody. Describe what happens to a rover's image file, a caution and warning message and a voice call when the relay drops below the horizon mid transfer. EXTEND

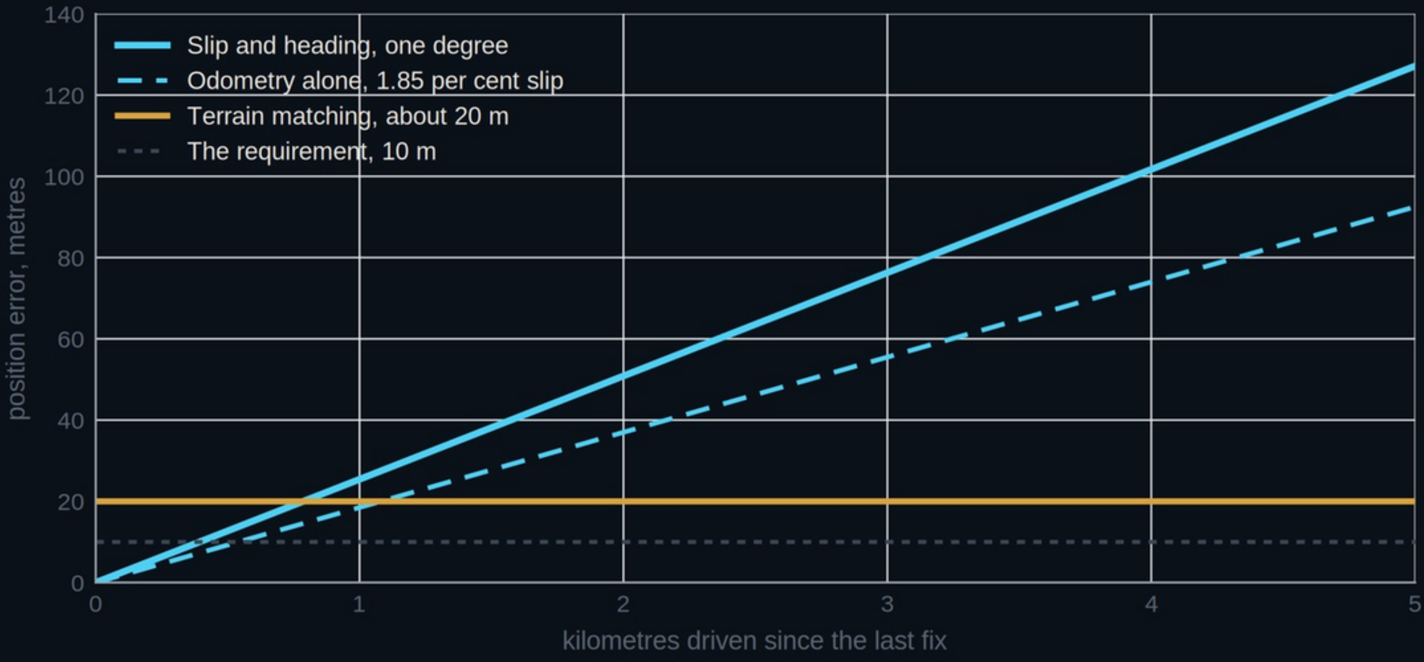

Navigation

No compass, no north, a Sun that circles the horizon and craters that all look the same. Every source of position on this site is wrong in its own way, so the design carries five of them and lets a filter argue between them.

WHAT THIS CHAPTER WILL TEACH YOU

1. Say why a compass, the Sun and the look of the ground are all useless here, and why the first of the three is worse than useless.
2. Separate a source that counts from a source that measures, and say how the error of each behaves as the distance grows.
3. Work out the position error of a dead reckoned traverse from wheel slip and heading error.
4. Say what accuracy the operations on this site actually demand, and which one of them sets the number.

10 m
ACCURACY REQUIRED

0 today
LUNAR NAVIGATION SERVICES

2.4 km
HORIZON AT EYE HEIGHT

69 m
SHADOW OFF A STANDING CREW MEMBER

15.1 Everything you would normally use

Take a compass out onto the surface and the needle will settle on something. It will not be north.

The Moon lost its global field billions of years ago, and what is left is crustal patches of tens to a few hundred nanotesla, the strongest mapped one 718, local and pointing whichever way they were frozen in. A compass here is worse than a broken instrument, because a broken instrument does not give you an answer.

The Sun is no better. At this site it circles the whole sky once a month and never climbs more than about a degree and a half, so there is no noon and no south. A crew member 1.8 metres tall throws a shadow 69 metres long, and that shadow swings through every bearing in the course of a month. Everything anybody has ever learned about reading direction from the Sun was learned under a different sky.

The ground is the third disappointment. Grey regolith under a hard sun with no air to soften distance, and every crater looking like every other crater. Judgement of size and range fails badly with no trees, no buildings and no haze to calibrate against, and Apollo crews said so at the time. So position on the Moon is not looked up. It is built, out of sources that are each wrong in their own way, and the trick is to pick five that are not wrong together.

> **69 metres** The shadow a standing crew member throws. The Sun here never rises more than about one and a half degrees, and a shadow is the height divided by the tangent of that angle. For somebody 1.8 metres tall it comes to 69 metres. The Sun at this site is a clock, not a compass.

Table 15.1. The five sources, and the different way each one lets you down.

Source	What it gives	How it fails	Where it works
Inertial unit	Motion, continuously, at any speed	Drifts, and no drift is ever recovered	Anywhere, for a few hundred metres
Odometry	Distance and heading from wheels and camera	Slips in regolith, 1.85 per cent assumed	Anywhere, worse with every metre
Four surveyed beacons	An absolute fix, about ten metres	Stops dead at the edge of the network	Inside the mast's 5.4 km horizon
Star tracker	Attitude, and the sky is black all day	Gives no position whatsoever	Anywhere, any time
Terrain matching	A fix of about twenty metres	Needs terrain distinctive enough to match	Anywhere the orbiter has imaged

Each row fails in a way no other row does, which is the reason for carrying all five. A Kalman filter weights each by how far it can be trusted at that moment, and the estimate is never worse than the best source available.

15.2 Counting against measuring

There are only two ways to know where you are. Count how far you have come from somewhere you knew, or measure where you are against something fixed. Every instrument in the table does one or the other, and the two fail in completely different shapes.

Counting is dead reckoning. An inertial unit feels the turns and the accelerations, the wheels count their own revolutions, and the camera watches the ground go past. It is always available and it gets worse with every metre. Wheels slip in regolith, and Apollo's rover analysis assumed 1.85 per cent: on that assumption Apollo 16 and 17 closed their traverses inside the 100 metre specification, but a climb up North Ray Ridge on Apollo 16 slipped enough to put the readout 250 metres out.

Measuring does not accumulate at all. Four surveyed beacons at the corners of the working area range to anything inside it by time of flight, over the same cellular network that carries the data. Beyond them a rover matches its camera view against the Lunar Reconnaissance Orbiter model, which resolves half a metre per pixel from 50 kilometres up. Craters all look alike to a person; to a matcher the pattern of their sizes is a fingerprint. Mars 2020 put Perseverance about 5 metres from the point its vision system picked, against a requirement of 60.

THE ONLY FORMULA IN THIS CHAPTER

$$E = s \times L \quad \text{and} \quad E = \theta \times L$$

E is the position error at the end of a leg in metres, s the wheel slip, the fraction of the wheels' turning that never becomes ground covered, L the distance driven since the last fix in metres, and θ the heading error built up since the last fix in radians. The slip error lies along the track and the heading error throws the machine sideways, so the two combine at right angles, by Pythagoras rather than by addition.

WORKED EXAMPLE 15.1 · HOW FAR WRONG A ROVER IS AFTER THREE KILOMETRES

Apollo's 1.85 per cent slip, a heading error of one degree built up since the last fix, a leg of 3 kilometres, and a terrain matched fix allocated at 20 metres that does not grow with distance driven.

1. Take the slip term first. It lies along the track, and the working is in metres.
   $E = 0.0185 \times 3\,000 \text{ m} = 55.5 \text{ m}$, call it 56 m
2. Now the heading term, which throws the rover sideways. One degree is 0.0175 radians.
   $E = 0.0175 \times 3\,000 \text{ m} = 52.5 \text{ m}$, call it 53 m
3. Combine them. They are at right angles to each other, so it is Pythagoras and not addition.
   $\sqrt{(56^2 + 53^2)} = \sqrt{5\,945} = 77 \text{ m}$

4. Find where counting stops beating measuring. Use the slip term alone, which is generous to the counting.

   L = 20 m ÷ 0.0185 = 1 081 m, about 1.1 km

ANSWER About 77 metres after three kilometres, against 20 metres from a camera that looks at the ground and asks where it is now. Past about 1.1 kilometres, counting loses to measuring and it never wins again, because one error grows with every metre and the other does not. That is why terrain matching is carried for everything outside the beacon network, why the inertial unit is calibrated against the beacons on every return to base, and why odometry is never trusted on its own. One caution about the 1.85 per cent: it is Apollo's assumption for mare and highland soil at low latitude with a crew aboard. Nobody has measured what a wheel does on polar regolith.

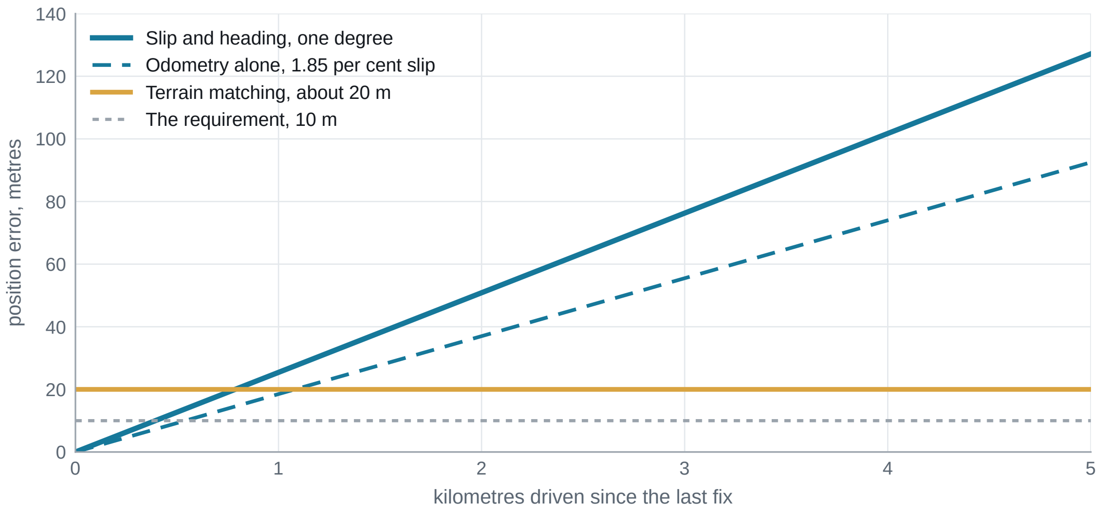

Figure 15.1. Position error against distance driven since the last fix. The sloping lines accumulate: at Apollo's 1.85 per cent slip, odometry alone reaches 93 metres after five kilometres, and one degree of heading error on top of it makes 127. The flat lines do not accumulate at all, because a terrain match and a beacon range each measure where the machine is now instead of adding up where it has been. The slip line crosses the terrain fix at about 1.1 kilometres and the combined line crosses it sooner, at about 790 metres. All four are design allocations for this concept and not measurements.

WHERE PEOPLE GO WRONG

GPS reaches the Moon, so surely this is solved. It does reach: on 3 March 2025 the LuGRE receiver on the Blue Ghost lander tracked GPS and Galileo signals at the Moon and computed a fix, the first ever obtained there. The catch is geometry. Every satellite of both constellations sits in almost the same direction from the Moon, a small patch of sky around Earth, so the fix is good along that line and poor across it, and a receiver at the pole sees the whole constellation on the horizon or below it. A signal designed to be used from inside its own constellation is a different instrument when used from four hundred thousand kilometres outside.

15.3 Ten metres, and who sets it

Ten metres is not a wish. Three operations set it. A lander has to arrive on a 22 metre apron, so terminal guidance needs better than ten to have any margin at all. A rover coming home through the dust it is throwing up itself has to find a two metre hatch. And a crew member walking back on a low battery has to be pointed at the base rather than past it, because with a 2.4 kilometre horizon at eye height, walking past your own habitat is fatal.

That last one is the one that decides it. The walk back rule is that nobody goes further out than they can walk home on their reserve, and that sum is only as good as the position estimate it starts from. Navigation on this base is a life support system under another name.

There is still no operational lunar navigation service. Technology demonstrators are flying, and NASA's first relay satellite is planned to fly in a network of five. ESA's Lunar Pathfinder has slipped to no earlier than 2027, and Moonlight's initial service is planned for the end of 2028 with full service by 2030. The LunaNet specification already defines an Augmented Forward Signal, a satellite navigation style broadcast a surface receiver could take a fix from. Until a spacecraft transmits one, every position on this site is measured from four beacons the crew surveyed themselves.

> **Why four beacons and not three.** Three ranges are enough to fix a position on a surface. The fourth is there for the day one of the others fails. With one of the four offline the geometry degrades and the accuracy falls to about six metres. That is acceptable, it is logged, and the beacon is repaired on the next excursion.

QUESTIONS

1. A rover drives 4.5 kilometres beyond the beacon network on odometry and gyro alone. Take Apollo's 1.85 per cent wheel slip and a heading error of half a degree, which is 0.00873 radians. Find the position error at the end of the drive, and say how many two metre hatch widths that is. APPLY

2. Explain why a magnetic compass on the Moon is not merely useless but misleading, and why a star tracker, which works perfectly there, still cannot tell a rover where it is. EXPLAIN

3. Every position on this site is measured from four beacons the crew surveyed themselves at installation. Say what that dependency costs, and what would have to happen before it could be given up. EXTEND

4. Show that the combined slip and heading line in Figure 15.1 crosses the 20 metre terrain fix at about 790 metres, and explain why that is sooner than the 1.1 kilometres of the slip line alone. CALCULATE

5. A crew member 1.8 metres tall throws a 69 metre shadow. Find the Sun's elevation at which the shadow would be 30 metres, and say whether that ever happens at this site. CALCULATE

6. The walk back rule limits every excursion to the distance a crew member can walk home on reserve. Say what quantities go into that sum, and which of them the navigation system is responsible for. EXTEND

complexity increases to the left, authority to move runs to the right, and the two never cross

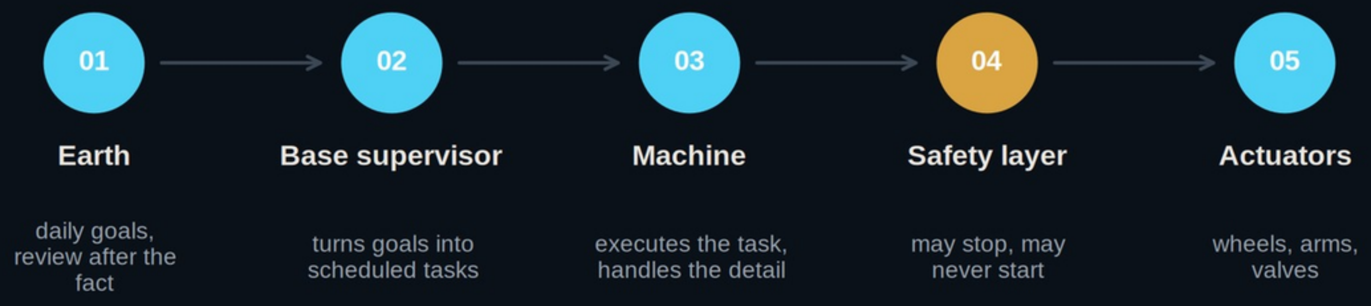

Earth never commands an actuator; the safety layer holds exactly one authority, to stop

Autonomy

Every command from Earth arrives late and every consequence is seen later still. That one fact decides where the decisions on this site are made, and it turns autonomy from a feature into a requirement.

WHAT THIS CHAPTER WILL TEACH YOU

1. Work out how far a machine travels blind between an operator seeing a hazard and the machine acting on it.
2. Place a task at the right level of autonomy, and defend the choice by its consequence rather than by its difficulty.
3. Describe how a fault is detected, isolated and recovered from with nobody in the loop.
4. Judge autonomy by the crew hours and the dose it saves rather than by the watts it costs.

2.56 seconds
ROUND TRIP TO EARTH

1.7 metres
BLIND TRAVEL PER COMMAND

94 per cent
MACHINE HOURS UNSUPERVISED, TARGET

138 watts
COMPUTE DRAW

16.1 Two and a half seconds is too long to drive anything

Earth is 1.28 light seconds away, so a command and the picture that answers it take 2.56 seconds to make the trip, and nothing anybody builds will shorten that half of the loop.

Add ground processing, the relay hops and the operator's own reaction, and 3.4 seconds is the honest figure. Now turn it into distance. A machine crossing the site at a walking pace of 0.5 metres a second travels 1.7 metres between the moment an operator sees something in a camera frame and the moment the machine acts. A boulder that appears at the edge of the frame is already under the wheel. Apollo sidestepped this by putting the driver aboard the rover. An uncrewed construction site has no such option.

The delay does not forbid remote driving. It prices it, and the price is speed. Studies of delayed teleoperation find task performance falling away between half a second and three seconds of lag, and this site sits at the far end of that range on its best day. So the decisions are made where the work is. The base plans its own day, each machine handles what happens inside its own task, and Earth keeps the two jobs a slow link is good for: setting goals and reading the logs. Nothing on Earth commands an actuator.

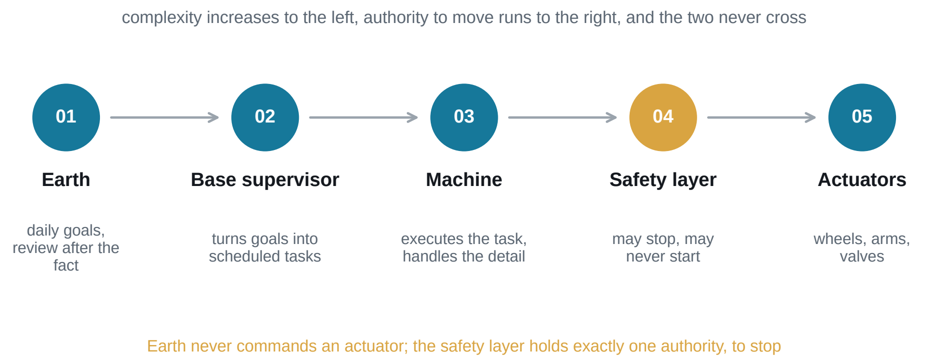

Figure 16.1. Where the decisions are made. Earth sets goals on a daily cadence and never commands an actuator. The supervisor turns goals into scheduled tasks, the machine handles everything inside a task, and the safety layer sits under all of it with one authority: to stop. The crew enter at the supervisor, for exceptions and for the work close to the habitat.

16.2 What may be delegated, and what may not

Table 16.1. The four levels this base uses, and what each one means on the ground.

Level	What it means	Where it is used	The crew's part
Full	The machine plans and acts inside a task without asking	Excavation, haulage on a graded route, laser sintering, running the plant	Set the day's goals, read the log
Supervised	The machine acts, a human watches and approves each step	Placing regolith over a module, work near the hulls	Approve each step, stop at any time
Teleoperated	A human commands the motion directly, at reduced speed	Recovery from a fault the machine cannot classify	Drive, slowly, from the base
Inhibited	The machine is prevented from moving at all	Whenever a crew member is inside its working envelope	Nothing, and that is the point

The design credits its levels to NASA's TX10 taxonomy, which lists capability areas and contains no levels of autonomy at all. The ladder this table follows is ECSS E ST 70 11C: E1 commanded from the ground, E2 a stored sequence run against the clock, E3 reacting to events on board, E4 planning towards a goal. The supervisor here runs at E4.

Read the third column and a pattern appears, and it is not the one most people expect. Excavation is the most repetitive work on the site and it is handed over completely. Lifting a tonne of regolith over the roof of a module is no harder to compute, and it stays under a human hand. The level is set by consequence, not by difficulty. A task is held close when getting it wrong cannot be undone, and released when the worst case is a wasted hour and a bucket to refill.

The design target is that 94 per cent of machine hours run unsupervised, and that number has no flight precedent anywhere. The longest autonomous drives on another world are Perseverance's: 347.7 metres in a single Martian day, 699.9 metres without human review, and a ground team planning every sol behind both. NASA's ISRU Pilot Excavator is built to dig and deliver 10 tonnes of regolith over 11 days on its own software, and it has not flown. So 94 per cent is a figure to be earned increment by increment, beginning with haulage on a graded road, and not a measurement.

WORKED EXAMPLE 16.1 · WHAT TELEOPERATION COSTS THE SCHEDULE

A loop of 3.4 seconds, a hauler that runs at 0.5 metres a second under its own control, a haul of 65 metres from the landing pad to the habitat, and a blind travel we are willing to accept of 0.1 metres. The 0.1 is our choice. Everything after it is arithmetic.

1. Find the blind travel at full speed.
   $d = 0.5\ \text{m/s} \times 3.4\ \text{s} = 1.7\ \text{m}$

2. Rearrange for the speed that keeps blind travel inside our tolerance.
   $v = 0.1\ \text{m} \div 3.4\ \text{s} = 0.029\ \text{m/s}$, about 2.9 cm/s

3. Time the 65 metre haul at that speed.
   $t = 65\ \text{m} \div 0.029\ \text{m/s} = 2\,241\ \text{s}$, which is 37 minutes

4. Time the same haul at the speed the machine uses when it is deciding for itself.

   t = 65 m ÷ 0.5 m/s = 130 s, about 2 minutes

ANSWER Teleoperation is not unsafe at 3.4 seconds of delay. It is roughly seventeen times slower, because 2,241 seconds against 130 is a factor of 17. One haul becomes the difference between a two minute trip and most of an hour, and this base has thousands of tonnes of regolith to move. That is the argument for autonomy in a line: the delay does not stop the work, it multiplies the calendar.

WHERE PEOPLE GO WRONG

The instinct is that a human in direct control is the safe arrangement and autonomy is the risk taken to save money. At this distance the instinct inverts. An operator on Earth is acting on a picture of the site as it stood 1.7 metres of travel ago. What makes a machine safe here is not a person at the controls but a safety layer on the machine itself, sitting below the autonomy, simple enough to be verified on its own and holding exactly one authority: to stop. The human is then put where judgement beats reaction time, which is setting the day's goals, approving a plan, and standing beside the work that happens close to the habitat.

16.3 Detect, isolate, recover, with nobody awake

A fault here has to be handled here. Detection is mostly disagreement: two sensors that should say the same thing and do not, a task that should have progressed and has not, a machine that should have answered its heartbeat within 5 seconds and is silent. Isolation is nearly always the same move, because stopping is safe from every state a machine can be in, and that is why the safety layer needs only the one verb. A deadlocked pair of machines is re sequenced by the supervisor. A machine that has lost its link stops, then drives itself home along the track it recorded on the way out.

One recovery is deliberately not automatic. The models are frozen for a week at a time, every decision they make is logged, and Earth reviews the logs before a new model is uplinked. A system that changes its own behaviour between reviews is a system nobody can certify, so the base gives up continuous learning in exchange for being answerable.

All of it runs on 138 watts of compute, 78 of them for perception, because seeing is the expensive part of deciding. Against the 27.6 kilowatt habitat load that is 138 ÷ 27,600 = 0.005, half of one per cent. The trade is not settled in watts anyway. The scarce thing here is the crew hour spent outside. An hour on the open surface costs 0.043 millisieverts of dose at the rate chapter 8 works from, and the Chang'e 4 dosimeter that rate is credited to read 1.37 millisieverts a day, which is 0.057 an hour, a third higher. Around each hour outside sit hours of suit preparation and dust carried back through the hatch. Four people have a decade of construction in front of them, and half of one per cent of the habitat load is what buys most of those hours back.

QUESTIONS

1 An operator drives a machine at 0.25 metres a second with a 3.4 second loop. How far does it travel blind, and what speed would hold the blind travel under 0.5 metres? APPLY

2 Explain why the safety layer runs on its own hardware, why its only authority is to stop, and why no layer above it may overrule a stop it has commanded. EXPLAIN

3 The 94 per cent unsupervised target has no flight precedent. Say what evidence collected on site would raise your confidence in it, and what would sink it. EXTEND

4 Table 16.1 hands excavation over completely and keeps placing regolith over a module under a human hand. Give one task on the site that you would move up a level and one you would move down, and defend each by its consequence. EXTEND

5 A rotation gives the crew about twenty working hours a day between them. If autonomy saves one crew hour outside for every four machine hours it runs, and the machines run 18 hours a day, estimate the crew hours and the dose saved in a 180 day rotation. CALCULATE

6 The models are frozen for a week at a time and Earth reviews the logs before a new one is uplinked. Say what that costs the base in adaptability, and why the design accepts the cost. EXPLAIN