Artificial Organs and Transhuman Trajectories

Physiological Redesign, Engineered Skin, and the Habitat-Carrying Body

Flyxion
Independent Researcher

November 2026

Preface

This book is organized as a graduate text in systems physiology and speculative engineering. Its first half supplies the tools needed to analyze an existing organ or a proposed replacement: a vocabulary of operations and boundaries, a way of stating what a person must remain able to do, conservation accounting for matter and energy, feedback analysis, scaling arguments, and reliability models. Its second half applies those tools to a single extended question, namely how much of the machinery that sustains a human life could migrate into the body, onto its surface, or into infrastructure that travels with it.

Two time horizons structure the later chapters. A horizon of one hundred years serves as a design horizon, within which a reader can trace plausible sequences of steps from technologies that already exist. A horizon of one thousand years serves as a comparative exercise, in which different bodies are examined under different constraints in order to see which constraints do the most work. Neither horizon is a forecast. Each is a device for forcing a design to state what it assumes.

The recurring question of the book is what has to remain possible for the person as the machinery that sustains them changes. The answer is stated in terms of physiological regulation, movement, sensation, repair, and the ability to survive an interruption in support. This formulation gives the claim about smaller bodies a precise form: whether the amount of material a person carries can decrease while the range of conditions in which that person can live and act stays the same or expands. Part V treats that claim as an optimization problem with explicit accounts, and it also identifies where the claim fails. Part VI turns to aging and repair, and supplies the test that the later chapters rely on: after an intervention, what can the person recover from that they could no longer recover from before?

Conventions

Every chapter closes its main exposition with a short evidence statement that separates three kinds of claim. Established physiology refers to findings that are standard in textbooks and replicated in the primary literature. Demonstrated engineering refers to devices or processes that have been built and operated, whether or not they have been used in a human body. Proposed design refers to concepts that exist as specifications, models, or manuscripts and have not been built. Several named designs recur throughout: the hepastitium, a distributed auxiliary organ; the Umbilicus, an elective adult exchange interface; continuous skeletal adjustment; micro-repair devices; the Fractal Brain Keel; and the integument that develops from a worn electric jacket. All of these belong to the third category, and the exercises repeatedly ask what would have to be measured to move any one of them into the second.

Quantities are given in SI units unless a physiological convention is much more common. Worked examples state their assumptions explicitly, and numerical values in them are illustrative parameter choices intended to fix orders of magnitude, not measurements of any particular person or device. Exercises are numbered within chapters. The bibliography appears at the end of the volume.

Contents

Preface **ii**

I The Concept of an Organ **1**

1 The Wet Animal **2**

1.1 Epithelia, Mucus, and the Meaning of Inside 2

1.2 Hard Material and Soft Material 2

1.3 The Body as an Ecosystem 3

2 Anatomy as an Allocation of Work **5**

2.1 Operations and Their Allocation 5

2.2 Five Kinds of Intervention 6

2.3 Operations Shared and Operations Spread 6

3 The Body and Its Boundaries **7**

3.1 Flux Across a Layered Boundary 7

3.2 A Boundary Is Active as Well as Passive 7

3.3 Selecting a Boundary for Analysis 8

4 Preserving Function **10**

4.1 Specifications as Admissible Trajectories under Causal Control 10

4.2 Reserve, Recovery, and Fallback 11

4.3 Specifications for Interruption 12

4.4 One Running Problem 12

4.5 What a Specification Cannot Say 13

II Accounting for Matter, Energy, and Control **14**

5 Conservation and the Metabolic Ledger **15**

5.1 The Balance Equation 15

5.2 Water, Gases, and Energy for a Reference Adult 15

5.3 Varying the Reference Person 15

5.4 Choosing the Boundary and Following the Dependence 17

5.5 Teaching the Habit of the Boundary 18

6 Separation, Sorting, and Water Recovery **19**

6.1 Separation by Density: What a Centrifuge Can Do 19

6.2 Separation of Dissolved Material 20

6.3 Three Versions of One Unit 21

7 Physiological Control Under Changing Demand **23**

7.1 Feedback, Delay, and Gain 23

7.2 Interacting Loops and Disagreeing Sensors 23

7.3 A Case of Competing Objectives 24

7.4 Design Consequences 25

8 Scaling, Redundancy, and Repair **26**

8.1 Geometric Scaling 26

8.2 Centralized and Distributed Units 27

8.3 The Hepastitial Threshold Model 27

8.4 Coverage and Connectivity 28

III Reassigning the Work of the Body **30**

9 Circulation, Respiration, and the Moving Body **31**

9.1 Oxygen Delivery and Extraction 31

9.2 Local Support and the Limits of Diffusion 32

9.3 How Close Must Support Be? The Krogh Cylinder 32

9.4 Interruption 33

10 Digestion, Renal Function, and Metabolic Exchange **34**

10.1 Regulation Is More Than Removal 34

10.2 Supplementing an Operation versus Assuming the Regulation 34

10.3 Intermittent Clearance as a Periodic Problem 35

10.4 Food, Heat, and the Timing of Energy 36

11 Structural Organs and Engineered Integument **37**

11.1 Rigidity, Load, and Scale 37

11.2 Remodeling and Its Rate Limits 37

11.3 Skin as an Organ and the External Alternative 38

12 Sensation, Actuation, and the Fractal Brain Keel **40**

12.1 Three Separable Descriptions 40

12.2 Cortical Surface Area and Its Constraints 41

12.3 The Fractal Brain Keel as a Proposal 41

IV Three Architectures for a Redesigned Body 43

13 The Hepastitium 44

13.1 Functions and Where They Live 44

13.2 Partial Support after Failure 45

13.3 Mesh, Coordination, and Common Causes 45

14 The Umbilicus 47

14.1 Three Modes of Coupling 47

14.2 What Must Remain in the Body 47

14.3 The Interface Itself 48

15 The Electric Jacket Becomes Skin 50

15.1 The Garment as a Boundary Layer 50

15.2 From Garment to Integument 50

15.3 Fluid Handling as the Distinctive Function 51

15.4 Regulation by Geometry: The Thermovariable Weave 51

16 The House, Car, Fridge, Stove, and Sink 53

16.1 Services and Their Scales 53

16.2 Sharing Hardware 53

16.3 The Design Problem 54

V Smaller Bodies, Larger Operating Ranges 55

17 Functional Compaction 56

17.1 Four Accounts 56

17.2 The Problem 56

17.3 Reading the Accounts 57

17.4 A Pareto Example 58

18 Where Miniaturization Stops Paying 60

18.1 Heat at Fixed Function 60

18.2 Reserves and Endurance 60

18.3 Shrinking the Component and Shrinking the System 61

18.4 Maintenance at Small Scale 61

19 Development Into a Different Anatomy **62**
19.1 Mechanics in Development 62
19.2 What the Framing Adds 62
19.3 Growth and the Renewal Schedule 63
19.4 Learning a New Location of Sensation and Control 63
19.5 Questions for a Developmental Specification 63

20 Maintenance, Dependence, and the Ability to Leave **65**
20.1 A Model of Eroding Native Capacity 65
20.2 Three Representations of the Loss of Options 66
20.3 Design Rules 66
20.4 Feasibility, Permission, and the Limits of Morphological Freedom 66

VI Aging, Repair, and the Recovery of Admissible Futures **69**

21 Computational Physiology and the Simulator **70**
21.1 What a Simulator Is For 70
21.2 Verification and Validation 70
21.3 Limits of a Simulator in the Design of Organs 71

22 Compensation and the Loss of Regulatory Loops **72**
22.1 Redundant Loops and Reserve 72
22.2 Why Compensation Hides Decline 73

23 Rejuvenation as Recovery of Admissible Futures **74**
23.1 Admissible Futures 74
23.2 Three Levels of Claim 74
23.3 The Test for an Artificial Organ 75

24 Turnover, Damage, and Retained Mechanical History **76**
24.1 Retention and Turnover 76
24.2 Damage Under Constrained Repair 77

25 Maintenance Without Memory: Cohort Turnover in Complex Systems **79**
25.1 Experience in a Growing Population 79
25.2 The Half-Life of a Rationale 80
25.3 Hardening Under Accumulation 80
25.4 Complexity as a By-Product of Incompatibility and of Standards 81
25.5 Games as a Contrast Case 81

25.6 Consequences for a Body That Carries Its Habitat 81

26 Restoration Before Explanation **83**

26.1 Examples 83

26.2 Three Tiers of Evidence 83

26.3 Design Consequences 83

27 Further Proposals in Architecture, Repair, and Cognition **85**

27.1 Forest-Inspired Organ Architecture 85

27.2 Tissue Repair as a Traversing Process 85

27.3 Cognitive Anti-Senescence 86

VII A Hundred to a Thousand Years **87**

28 The First Century **88**

28.1 Three Services 88

28.2 Stall Points as Results 89

29 The Next Several Centuries **90**

29.1 Why Lineages Persist Together 90

29.2 Disturbances 90

30 The Thousand-Year Body **92**

30.1 The Accounting Exercise 92

30.2 Operating Range 92

30.3 What the Comparison Supports 93

31 Capstone on the Minimum Habitable Body **94**

31.1 The Specification 94

31.2 The Model 94

31.3 Comparison and Transition 95

31.4 Criteria 95

31.5 What the Book Has Argued 95

A Transport and Separation **96**

A.1 Series Resistance 96

A.2 Minimum Work of Separation 96

A.3 Settling and Sedimentation Equilibrium 96

A.4 Heat Conduction from a Sphere 96

B Feedback and Stability **98**

B.1 Loop with Delay 98

B.2 Saturating Actuators 98

B.3 Competing Controllers 98

C Reliability Models **99**

C.1 The k-out-of-n Model 99

C.2 Heterogeneous Quality 99

C.3 Common Cause 99

C.4 Marginal Node 99

C.5 Coverage by a Poisson Process 100

D The Mathematics of Feasible Operating Ranges **101**

D.1 Viability 101

D.2 Grace Period 101

D.3 Intersections and Pareto Sets 101

D.4 Sensitivity of the Range to a Parameter 101

Bibliography **103**

Part I

The Concept of an Organ

Chapter 1

The Wet Animal

A human body is a column of water held in shape by a small quantity of mineral and protein, and every question about replacing or supplementing an organ eventually becomes a question about how that water is contained, moved, and conditioned. Water makes up roughly sixty percent of the mass of an adult male body and somewhat less of an adult female body, with the difference accounted for mainly by body composition. About two thirds of that water lies inside cells, and the remainder is extracellular, divided between the plasma, which circulates, and the interstitial fluid, which bathes the cells and moves only slowly. The interstitial compartment is the working environment of most tissue: oxygen, nutrients, signaling molecules, and waste products pass through it on their way between capillaries and cells, and it is the compartment in which any implanted device that is not in the bloodstream will sit.

1.1 Epithelia, Mucus, and the Meaning of Inside

The body has far less true interior than its anatomy diagrams suggest. The lumen of the gut, the airways, the bladder, and the ducts of glands are continuous with the outside world, and the tissue that lines them is a barrier that is continuously renewed and continuously selective. An epithelium is a sheet of cells joined by tight junctions, with distinct apical and basolateral faces that carry different transport proteins. This asymmetry allows directional transport: a renal tubule cell can pick up sodium from the lumen across one face and pump it out across the other, so that salt and, by osmosis, water move from the filtrate to the blood without any pipe or valve.

Mucus is the second element of the boundary. It is a hydrogel of heavily glycosylated proteins called mucins, and it traps particles, lubricates movement, and supports a layered chemical environment in which the oxygen, pH, and nutrient concentrations differ from those of the tissue beneath. In the airways the mucus layer is cleared by cilia on a timescale of tens of minutes, and in the intestine it is renewed on a timescale of hours. A mucosal surface is therefore a place where material is selected, conditioned, and turned over, and any proposed artificial interface that crosses such a surface has to be analyzed as a perturbation to a living selective system and not as a hole in a wall.

1.2 Hard Material and Soft Material

Animals differ in where they put mineralized tissue. Vertebrates place bone inside, surrounded by soft tissue, so that muscle pulls on levers that are shielded from the environment and the skeleton can grow and remodel throughout life. Mollusks with shells carry a mineralized enclosure outside a soft body, which gives protection that is adjusted slowly

by deposition at the margin. Arthropods wear a jointed exoskeleton of chitin and protein, sometimes mineralized, and must shed it periodically to grow, passing through a period of vulnerability after each molt. These solutions allocate the same functions of support, protection, and leverage to different places, and each allocation has a signature cost. An internal skeleton is metabolically expensive to maintain but permits continuous growth. An external skeleton is cheap to maintain in the intervals between molts but couples growth to a crisis. The comparison is useful later because the proposals examined in this book move structure toward the surface of the human body, and the arthropod case is a reminder that a hard exterior is entangled with the processes of growth and repair.

Muscle is the generator of movement in all three cases, and it is a large consumer of the body's energy budget during exertion. Skeletal muscle makes up roughly forty percent of lean body mass, and the heat released by contracting muscle is the largest variable term in the human thermal balance. This is why movement, heat, and water are not separable topics in the chapters that follow.

1.3 The Body as an Ecosystem

The human body carries a resident population of microorganisms, concentrated in the colon, that is comparable in number to the body's own cells. A widely cited estimate puts the ratio of bacterial to human cells near one to one for a reference adult, with a bacterial count on the order of 4×10^{13} [28]. These organisms transform material that the host cannot digest, produce vitamins and short-chain fatty acids, and compete with pathogens. The practical consequence for engineering is that the chemical exchanges in the body are shared among the host tissues and a resident biota, and that a device which changes the chemistry of a compartment, for example by removing water from the colon or altering the delivery of nutrients, changes the conditions of a second ecosystem as well.

It is useful to separate two meanings of chemical activity. The first is the repertoire of reactions, the set of distinct transformations of which the body is capable. Genome-scale reconstructions of human metabolism catalog several thousand distinct reactions. The second is the number of reaction events occurring per unit time throughout the whole body, which depends on how many cells, molecules, and organisms carry out those reactions and how fast. These are independent quantities. A system can have a small repertoire executed enormously many times, as in a bulk chemical plant, or a large repertoire executed rarely. The body is distinguished by combining a large repertoire with a very large event rate, and the event rate is what sets the heat output and the water and oxygen demand.

Worked Example 1.1 (Reaction events in a day)**.** A resting adult dissipates about 100 W, or 8.6 MJ per day. Take the free energy released by hydrolysis of one mole of ATP under cellular conditions to be about 55 $\mathrm{kJ\,mol^{-1}}$ and assume that all of the metabolic power is routed through ATP. The molar turnover is then

$$n_{\mathrm{ATP}} \approx \frac{8.6 \times 10^{6}\ \mathrm{J}}{5.5 \times 10^{4}\ \mathrm{J\,mol^{-1}}} \approx 160\ \mathrm{mol\ per\ day}.$$

With a molar mass of 0.507 $\mathrm{kg\,mol^{-1}}$ for ATP, this corresponds to roughly 80 kg of ATP cycled through its hydrolysis and resynthesis, close to the mass of the body itself, even though the body contains only tens of grams at any instant. The number of hydrolysis

events is about $160 \times 6.0 \times 10^{23} \approx 10^{26}$ per day, or about 10^{21} per second. The same repertoire of reactions is being executed at this rate throughout the body, and the figure illustrates why the event rate, and not the size of the repertoire, governs the water, oxygen, and heat budgets taken up in Part II.

Evidence Status

Established physiology. The composition of body water, the structure of epithelia and mucus, the position of mineralized tissue in vertebrates, mollusks, and arthropods, and the scale of the resident microbial population are textbook material [5, 18]. Demonstrated engineering. None of the devices discussed later is introduced in this chapter. Proposed design. The chapter introduces only the viewpoint on which the proposals depend, namely that a body is a boundary-managed volume of water and not a set of separate machines.

Exercise 1.1. Using a body mass of 70 kg and a water fraction of 0.6, estimate the volume of each fluid compartment with the fractions given in the text. Explain which compartment an implanted sensor that is not in contact with blood would sample, and why the answer affects the delay between a change in blood composition and the sensor reading.

Exercise 1.2. Repeat Worked Example 1.1 for a person sustaining 400 W of metabolic power during exertion. State which quantities change and which do not, and identify the one that bears most directly on the demand for oxygen delivery.

Chapter 2

Anatomy as an Allocation of Work

Anatomical names group tissue by position and appearance, whereas engineering names group it by what it does. Reading a body as an allocation of work to structures requires a vocabulary of operations, and the choice of that vocabulary determines which replacements and supplements can be described at all. A pump, a sorter, a store, an actuator, a signaling channel, and a repair process are operations.

2.1 Operations and Their Allocation

An operation is a transformation of one or more input streams into output streams under a stated control. The operations that recur across the body are sensing, pumping, sorting, storing, contracting, signaling, and repairing. Sensing converts a physical or chemical quantity into a signal. Pumping adds pressure or head to a fluid. Sorting separates a mixture into streams of different composition. Storing holds material or energy for later release. Contracting converts chemical energy to mechanical work. Signaling carries information between locations. Repairing restores a structure to a prior state or to a functional equivalent.

Definition 2.1 (Allocation matrix)**.** Let $O = \{o_1, \dots, o_m\}$ be a set of operations and $A = \{a_1, \dots, a_n\}$ a set of anatomical structures. The allocation matrix is the nonnegative array $C = (c_{ij})$ in which c_{ij} is the capacity with which structure a_j performs operation o_i, expressed in the units appropriate to the operation, such as volume flow for pumping or mass flux of a given solute for sorting.

The row sum $\sum_j c_{ij}$ is the total capacity of the body for operation o_i, and the number of nonzero entries in a row measures how widely the operation is distributed. The kidney, for example, sorts water and solutes, stores nothing of consequence, signals through the secretion of renin and erythropoietin, senses oxygen delivery, and contributes to repair through the activation of vitamin D, so its column has several entries of very different size.

Worked Example 2.1 (Reading the kidney as a column)**.** Healthy kidneys filter plasma at a glomerular filtration rate of about 180 L per day, and reabsorb all but about 1.5 L of it. The sorting operation thus handles a flux of roughly 125 $\mathrm{mL\,min^{-1}}$ of ultrafiltrate, of which about 99 percent is returned. Dialysis in its conventional form performs a portion of this sorting for solutes such as urea and potassium and for excess water, and it does so intermittently, typically for a few hours several times a week. It does not perform the endocrine operations. In allocation terms, dialysis supplies part of one entry of the kidney's column on an intermittent schedule, and a transplanted kidney supplies the whole column. The two are different kinds of intervention even though both are called renal replacement.

2.2 Five Kinds of Intervention

In replacement, a structure's whole column is supplied by a substitute, as when a cochlear implant takes over the transduction that hair cells carried out. In supplementation, part of a capacity is added to a structure that continues to operate, as with a pacemaker that paces when the intrinsic rhythm fails, or a ventricular assist pump that adds flow to a weakened heart. In externalization, an operation is carried out outside the body with the body's fluids routed to it, as in hemodialysis and in extracorporeal oxygenation. In redistribution, an operation formerly concentrated in one structure is spread over many, which is the central idea of the hepastitium treated in Chapter 13. In addition, a function is created that the original anatomy never performed, such as magnetic sensing or the continuous recovery of water from fluids that would otherwise be lost.

Each kind has a characteristic failure. Replacement fails by omission of entries in the column that no one thought to list. Supplementation fails when the native structure and the device compete or interfere. Externalization fails when the connection fails, so its reliability is limited by the interface. Redistribution fails at the points where coordination is needed. Addition fails if the new function interacts with regulation that the original designers of the system, namely evolution, tuned without it.

2.3 Operations Shared and Operations Spread

Two relations between operations and structures deserve names. An operation is concentrated when one structure carries nearly all of its capacity and spread when many do. Pumping of blood is concentrated in the heart, although skeletal muscle contraction contributes to venous return. Heat rejection is spread over the skin, airways, and circulation. Concentration gives simplicity and a clear target for replacement, and spreading gives tolerance to partial failure. A structure that performs several operations is multiplexed, and a failure in it removes several rows at once. The design of auxiliary organs can be viewed as the choice, for each operation, of how concentrated to make it and how many other operations to place on the same hardware.

Evidence Status

Established physiology. The functions of the major organs and the quantities used in Worked Example 2.1 are standard [5]. Demonstrated engineering. Replacement, supplementation, and externalization are all in clinical use, represented by cochlear implants, pacemakers, ventricular assist devices, and dialysis [22]. Proposed design. Redistribution and addition appear in this book mainly as proposals, and each is revisited in Parts IV and VI.

Exercise 2.1. Construct a six-by-six allocation matrix with the operations listed in this chapter and six structures of your choice. Use qualitative entries of zero, small, and large. Identify one operation that is concentrated, one that is spread, and one structure that is multiplexed.

Exercise 2.2. For each of the five kinds of intervention, give one existing clinical example, and name the entry of the allocation matrix that it changes. Where no clinical example exists, explain why.

Chapter 3

The Body and Its Boundaries

A boundary is a place where exchanges are selected and regulated. A material that passes from the environment into the circulation crosses several boundaries in sequence, each with its own selectivity and its own response to conditions, and the amount that arrives is governed by all of them together. The preceding chapters treated the skin and mucosa as edges of the body. This chapter treats them as one layer among several, and it introduces the series model of transport that will be used to place clothing and engineered integument on the same footing as tissue.

3.1 Flux Across a Layered Boundary

Let a solute or a vapor cross a layer of thickness ℓ and cross-sectional area A under a driving difference Δc in concentration, or in the appropriate potential. Over a wide range of conditions the steady flux follows a linear law,

$$J = \frac{\Delta c}{R}, \qquad R = \frac{\ell}{D\,\phi}, \tag{3.1}$$

where D is a diffusivity, ϕ a partition or porosity factor, and R the resistance per unit area. When several layers lie in series, the steady flux through each is the same, and the resistances add,

$$J = \frac{c_{\mathrm{in}} - c_{\mathrm{out}}}{\sum_k R_k}. \tag{3.2}$$

The layer with the largest resistance dominates. For water vapor leaving the body, the layers include the skin, a still layer of air trapped within or beneath clothing, the fabric itself, and the boundary layer of moving air at the outer surface. For heat the same structure holds with thermal resistances, which are conventionally given in units of $\mathrm{m^2\,K\,W^{-1}}$ and in clothing practice as the unit *clo*.

3.2 A Boundary Is Active as Well as Passive

Equation (3.2) describes passive transport, and the body's boundaries are not passive. Sweat glands change the vapor source term on demand, vasomotor control changes the heat delivered to the skin by an order of magnitude, and the epithelia of the gut and kidney pump solutes against gradients. The resistance of a layer in the body is therefore a controlled quantity and not a constant. A garment can add a resistance in series, but the body responds by changing its own, and the net effect of an added layer is the result of that interaction. This observation is the seed of the analysis of control in Chapter 7.

Worked Example 3.1 (One parcel of water, three boundaries)**.** Consider a parcel of water of one liter that enters the body as a drink, is excreted in sweat after passing through the

circulation, and is then either lost, retained in a garment, or recovered by a stationary system. The accounting depends entirely on where the boundary is drawn.

If the boundary is the skin, the parcel enters at the gut, is a part of the body's water for some hours, and leaves at the skin. Its residence time is set by the body's water turnover, which is about six percent of total body water per day at rest, giving a mean residence time of roughly two weeks for the body as a whole. A parcel that is routed to the sweat glands during exertion can leave within hours, so the residence time of any particular parcel depends on what the person does. The accounting shows a net loss of one liter at the skin.

If the boundary is a worn suit that encloses the skin and does not let liquid escape, the parcel leaves the body but not the system. It is now stored in the garment, the suit's mass has risen by one kilogram, and the thermal consequences of its evaporation are shifted to wherever the suit rejects heat. The accounting shows a transfer from one internal store to another, and the net loss is zero.

If the boundary is a larger system that includes a shared maintenance facility, the parcel leaves the suit at a scheduled service event, is treated, and returns to a supply. In this accounting, the quantity that is lost is only the fraction that the facility cannot recover, and the accounting must include the energy and consumables of the treatment. A loop that appears closed at the second boundary is open at the third, because recovery depends on inputs from outside it.

The question of whether a loop is closed has no answer until a boundary is chosen, and the choice of boundary is a design decision with consequences for who bears the cost of maintaining the loop.

3.3 Selecting a Boundary for Analysis

Three rules are useful in practice. The boundary should be drawn so that every flow crossing it can be measured or specified, which excludes boundaries that cut through a well-mixed compartment. It should be drawn so that the control actions of interest are inside it, so that the analysis can include the feedback that determines the resistance. And it should be drawn so that the maintenance of the system appears as a flow across it, which makes dependence on consumables, power, and service visible. A boundary that satisfies these rules will often not coincide with the skin.

Evidence Status

Established physiology. Equations (3.1) and (3.2) are standard transport relations, and the physiology of sweat, vasomotor control, and epithelial transport is textbook material [4, 18]. Demonstrated engineering. Vapor-permeable and liquid-handling garments exist, as do liquid-cooled garments for protective work and spaceflight [7]. Proposed design. The shared maintenance system in Worked Example 3.1 is a schematic of the later proposals and not a specification of any existing facility.

Exercise 3.1. A vapor-permeable fabric has a resistance to water vapor transport that is three times that of the still air layer beneath it, which in turn is equal to that of the outer boundary layer. Using Equation (3.2), find the fraction of the total resistance contributed by the fabric. If the outer boundary layer resistance is halved by wind, find the new fraction

and the factor by which the evaporative flux rises for a fixed driving difference.

Exercise 3.2. Redraw the boundary in Worked Example 3.1 so that it encloses only the shared maintenance system and the suit, excluding the body. State what is then counted as an input, and discuss whether the loop is closed.

Exercise 3.3. Three layers in series have vapor resistances of 2, 6, and 2 in the same units. Find the fraction of the total contributed by each layer. Then find the total resistance if the largest layer is halved, and compare the improvement with the one obtained by halving each of the two smaller layers.

Exercise 3.4. The skin is an active layer whose resistance to heat flow can be reduced by a factor of ten by vasodilation. Suppose that, when constricted, it has a thermal resistance equal to one third of the total resistance of the garment and air layers in series with it. Find the total resistance in the constricted and dilated states, and the largest improvement in total resistance that control of the skin alone can provide.

Exercise 3.5. A body contains 42 L of water and turns over 6 percent of it per day at rest. Find the mean residence time of a parcel. If a day of exertion adds 4 L of sweat loss that is replaced by drinking, find the turnover fraction for that day and the mean residence time if every such day were the same.

Exercise 3.6. A garment encloses the skin and holds all of the sweat that reaches it. For each of three boundaries, the skin, the garment, and the garment together with a laundering facility, state whether the water loop appears closed, what crosses the boundary, and who bears the cost of the energy used to dry the garment.

Chapter 4

Preserving Function

An artificial organ is judged by whether the person who relies on it can still do what they could do before, under the conditions in which they could do it. The statement of that requirement is a functional specification. Writing one precisely is harder than it appears, because a device that matches average performance can fail a specification that is written in terms of conditions, and because the specification has to include what happens when support is interrupted.

4.1 Specifications as Admissible Trajectories under Causal Control

Let $x(t) \in \mathbb{R}^n$ be the physiological state of a person, including quantities such as core temperature, body water, plasma osmolality, oxygen saturation, and the state of any device. An architecture A is given by four things: its dynamics $\dot{x} = f_A(x, u, d)$, in which u is the action taken and d a disturbance; an observation map $y = h_A(x)$ that says what the person or device can sense; a set U_A of available actions; and resource constraints, on energy and consumables, that restrict which actions can be sustained. A policy is a rule that chooses the action at each time from the history of observations up to that time, and Π_A is the set of admissible policies. Policies are causal: a policy does not see a disturbance except through what the observation map shows of its effects.

Let $\mathcal{K}$ be the viable set, given by explicit measurable limits, and let $\mathcal{K}_0 \subset \mathcal{K}$ be a nominal region to which the person is expected to return. Let X_0 be a set of initial states, let T be a duration, and let a task be stated that the person must remain able to perform. An environment e comes with a class $\mathcal{D}(e)$ of disturbances, such as an exertion, a thermal condition, or the withdrawal of a support.

Definition 4.1 (Functional specification and operating range)**.** A specification is a quadruple $\mathcal{S} = (\mathcal{K}, X_0, T, \text{task})$. Architecture A passes the specification in environment e if

$$\forall x_0 \in X_0 \;\; \exists \pi \in \Pi_A \;\; \forall d \in \mathcal{D}(e) :$$
$$x_{x_0}^{A,\pi,d}(t) \in \mathcal{K} \text{ for all } 0 \le t \le T \text{ and the task is performed.}$$

The operating range $R(A)$ is the set of environments in which A passes. If the architecture observes only part of the initial state, the quantifier over x_0 is replaced by a single policy for all initial states consistent with the initial observation.

The order of the quantifiers carries the meaning. The policy is chosen after the initial state and before the disturbance, and it cannot depend on the disturbance except through observations. If the existential quantifier were placed after the universal one over d, each disturbance would receive its own controller that knows what is coming, and an architecture

would appear more capable than it is. Every later calculation in the book, including the operating ranges of Chapter 30, refers to this definition. The notation $\mathcal{D}$ without an argument denotes the disturbance class of the environments under consideration.

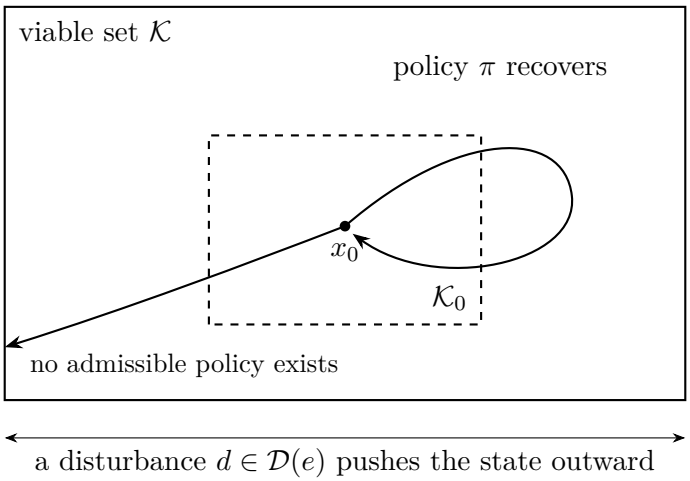


Figure 4.1: The state plane of Definition 4.1. The policy is chosen after the initial state x_0 and acts on what it observes. The right-hand trajectory is one for which a policy exists that keeps the state in $\mathcal{K}$ for the whole duration and returns it to $\mathcal{K}_0$. The lower-left trajectory leaves $\mathcal{K}$, and it stands for a disturbance for which no such policy exists, so the architecture does not pass the specification in that environment.

This formulation draws on the theory of viability, in which the set of initial states from which a controlled system can remain in a constraint set for all time is called the viability kernel [2]. The limits that define $\mathcal{K}$ must be given numerically in every worked example, and the task must be a stated activity. The phrase conscious, injury-free, and able to act motivates the choice and is not a substitute for it.

4.2 Reserve, Recovery, and Fallback

Four quantities refine the specification. The operating range is the set of environmental and activity conditions over which the specification holds. The reserve capacity ratio is

$$\rho_{\mathrm{res}} = \frac{C_{\max}}{D_{\max}}, \tag{4.1}$$

the ratio of the maximum capacity of a function to the maximum demand placed on it by $\mathcal{D}$. A ratio near one means that any additional disturbance exhausts the function. The recovery time τ_{rec} is the time required for the state to return to a nominal region after a disturbance ends. A fallback state is a configuration of reduced capability to which the person or system can move so that the state remains in $\mathcal{K}$ while support is unavailable. The grace period τ_g is the longest time for which the state remains in $\mathcal{K}$ after a support is withdrawn, given the best available fallback. These quantities appear in every later chapter, and Chapter 20 defines a graded notion of the ability to discontinue support in terms of τ_g and recovery. If the support is withdrawn with the person already carrying some accumulated damage D_0, the headroom is reduced accordingly, and Chapter 20 states the grace period with that

term included. Chapter 23 adds a test of what the person can recover from, which applies the specification to the question of whether a function has been restored.

Worked Example 4.1 (A pump that meets the mean)**.** A device is proposed to supply the entire circulatory flow of a person at a fixed output of 7 $\mathrm{L\,min^{-1}}$, chosen to match the person's time-averaged demand over a typical day. Resting cardiac output is about 5 $\mathrm{L\,min^{-1}}$, and the output of a healthy adult in intense exertion is typically 20 $\mathrm{L\,min^{-1}}$ or more. Suppose that the person spends 22 hours of the day at demands between 5 and 7 liters per minute and two hours at a demand of 15.

The time-averaged demand, taking a representative 6 $\mathrm{L\,min^{-1}}$ for the first 22 hours, is $(22\times 6+2\times 15)/24 = 6.75\ \mathrm{L\,min^{-1}}$, which is close to the pump's output, so the pump matches the mean. The reserve ratio from Equation (4.1), taking $D_{\max} = 15$, is $\rho_{\mathrm{res}} = 7/15 \approx 0.47$, which is below one, so the specification fails during the two hours of exertion despite the match on average. Fixed output also removes the ability to supply less than 7 $\mathrm{L\,min^{-1}}$ during sleep, when demand falls to about 4, and the surplus flow must be handled by changes in vascular resistance. The pump meets the mean and fails the specification, and the failure appears only because the specification was written in terms of conditions.

The example shows why modern ventricular assist devices are built to vary output with demand and with the native contribution of the heart, rather than to deliver a fixed flow [22].

4.3 Specifications for Interruption

A specification that is satisfied only while supply is uninterrupted is a specification about the supplier as well as the person. The requirement of survival through interruption is stated by adding to $\mathcal{D}$ the class of disturbances in which a support of the design, whether power, consumable, or connection, is lost for a stated duration. A design then has a grace period under each such disturbance, and the specification fixes a minimum. Chapters 14 and 20 use this machinery to compare designs that carry their support with designs that draw it through a connection.

4.4 One Running Problem

A single problem is carried through the book so that each new tool changes an answer that the reader already knows. The person is the reference adult of Chapter 5, at 70 kg and 175 cm, starting each day normothermic and euhydrated with no accumulated damage, so that X_0 is a single state. The viable set is a core temperature between 35.0 and 39.0 °C and a body water deficit of no more than 2.1 kg, three percent of body mass. The task is two hours of walking at a metabolic power of 300 W with 75 W of external work, within a day of duration $T = 24$ h. Three environments are used: e_1, an air temperature of 20 °C; e_2, an air temperature of 35 °C; and e_3, the conditions of e_1 with every external support withdrawn for three days. Three architectures are compared throughout: a worn garment G that recovers sweat, a hepastitial subsystem H that recovers fluid inside the body, and a service U connected through the Umbilicus. The question is which of them passes the

specification in which environments. Each return to the problem adds one constraint that can change the answer.

4.5 What a Specification Cannot Say

A viability specification says what must remain possible and not what is desirable. It does not distinguish between a person who can walk, work, and travel and a person who is merely kept inside $\mathcal{K}$. Additional constraints may be written to require particular activities in $\mathcal{D}$, and each added constraint narrows the set of admissible designs. The choice of specification is a value judgment that is made by the person who relies on the device, and the exercises that close each chapter in the later parts repeatedly return the question of what would have to be measured to establish that a design improves the life of that person.

Evidence Status

Established physiology. Resting and maximal cardiac output, and the dependence of demand on activity, are standard values [18]. Demonstrated engineering. Ventricular assist devices with speed modulation and magnetic levitation have been operated clinically [22]. Viability theory is a mathematical framework with established results [2]. Proposed design. The grace period and the interruption class $\mathcal{D}$ are used here as design quantities, and standard clinical practice does not report them as such.

Exercise 4.1. Construct a demand profile for one day, in the form of a table of activities, durations, and cardiac outputs, and compute the time average. Find the smallest constant pump output that meets the specification of Worked Example 4.1 for all activities, and compare the energy delivered by that pump with the time-averaged demand.

Exercise 4.2. A water recovery unit carried in a garment fails at the start of a day of exertion in a hot climate. Define the viable set $\mathcal{K}$ in terms of core temperature and body water deficit, give plausible bounds, and describe what a fallback state would be. State what must be measured to estimate τ_g for this case.

Exercise 4.3. Specify in words a function that a person could retain according to the viability definition above and still have lost something that matters to them. Propose an additional constraint that would exclude the outcome.

Part II

Accounting for Matter, Energy, and Control

Chapter 5

Conservation and the Metabolic Ledger

Every proposal in this book can be reduced to a statement about what crosses a boundary and what accumulates inside it. The ledger is the bookkeeping device that makes the statement checkable. It is deliberately plain, because its value lies in the discipline of choosing a boundary, listing every stream that crosses it, and refusing to let a quantity appear or disappear without a source or a sink.

5.1 The Balance Equation

For any conserved quantity M inside a chosen boundary, the balance is

$$\frac{\mathrm{d}M}{\mathrm{d}t} = \sum_{k} \dot{m}_k^{\mathrm{in}} - \sum_{k} \dot{m}_k^{\mathrm{out}} + \dot{g}, \tag{5.1}$$

where the sums run over the streams crossing the boundary and $\dot{g}$ is the net rate of generation by reaction inside it. For chemical elements and for total mass the generation term is zero. For a particular molecular species, such as urea or carbon dioxide, it is nonzero and is supplied by metabolism. For energy, the balance has the form of the first law, and metabolic heat release appears as a generation term inside the body boundary and as a chemical enthalpy inflow when the boundary encloses the food supply.

5.2 Water, Gases, and Energy for a Reference Adult

The reference quantities in Table 5.1 are the working values for the rest of the book. They describe a 70 kg adult at rest in a temperate environment and are rounded to the precision at which they are reliable. The water routes of intake and loss each total about 2.6 L per day, and the individual routes vary widely between persons and conditions.

The oxygen figure follows from a resting consumption of about 0.25 $\mathrm{L\,min^{-1}}$ at standard conditions, which is 11 $\mathrm{mmol\,min^{-1}}$ and about 16 mol per day, or 0.5 kg. The carbon dioxide figure follows from a respiratory quotient near 0.8. Thus the mass of gas exchanged with the environment is larger than the mass of water that is drunk, a fact that is easy to overlook when attention is directed at fluids. The energy released per liter of oxygen consumed is close to 20 kJ over a wide range of diets, which is the basis of indirect calorimetry and links the oxygen and heat entries of the table.

5.3 Varying the Reference Person

The reference adult is a convenience, and every number in the ledger changes with the person. Body surface area, which governs exchange with the environment, can be estimated from

Table 5.1: Daily ledger for a reference adult at rest in a temperate environment.

Quantity	Route	Value per day
Water in	Drinks	1.5 L
	Water in food	0.8 L
	Metabolic water	0.3 L
Water out	Urine	1.5 L
	Skin, insensible	0.5 L
	Respiratory	0.4 L
	Sweat, thermal	0.1 L
	Feces	0.1 L
Oxygen consumed	Respiration	0.5 kg
Carbon dioxide produced	Respiration	0.6 kg
Energy in food	Chemical	8.6 MJ
Heat released	Dissipation	≈ 8.6 MJ
Solute load in urine	Osmoles	≈ 600 mOsm

mass W in kilograms and height H in centimeters by the formula $A = 0.007184\, W^{0.425} H^{0.725}$ in square meters. Metabolic power scales approximately as mass to the three-quarter power, and total water turnover roughly follows it.

Worked Example 5.1 (Five people, one ledger)**.** Table 5.2 applies these relations to five body sizes, taking 100 W and 2.6 L per day for 70 kg and scaling both as $(W/70)^{3/4}$. The surface area of the 70 kg person at 175 cm is 1.85 m^2, and the value of 1.8 m^2 used in earlier examples is a rounded figure. Power per unit area varies by only about ten percent across the table, whereas area per unit mass varies by half: 0.026 $\mathrm{m}^2\,\mathrm{kg}^{-1}$ for the reference adult and 0.040 $\mathrm{m}^2\,\mathrm{kg}^{-1}$ for the child. A child therefore has relatively more surface through which to lose heat and water, and a heat capacity per unit area of about 87 $\mathrm{kJ\,m^{-2}\,K^{-1}}$ against 131 for the adult, so the same flux per unit area changes core temperature about half as fast again in the child. The crude scaling also understates the water turnover of children, whose fractional turnover per kilogram is larger, and it ignores sex, age, composition, and acclimatization. Its purpose is to show which entries of a design calculation are sensitive to the person, and a design that is checked only for the reference adult has not been checked.

Table 5.2: The ledger scaled to five body sizes.

Person	Mass (kg)	Height (cm)	Area (m^2)	Power (W)	Water (L/day)
Child	20	115	0.80	39	1.0
Small adult	50	160	1.50	78	2.0
Older adult	65	165	1.72	95	2.5
Reference adult	70	175	1.85	100	2.6
Large adult	90	185	2.14	121	3.1

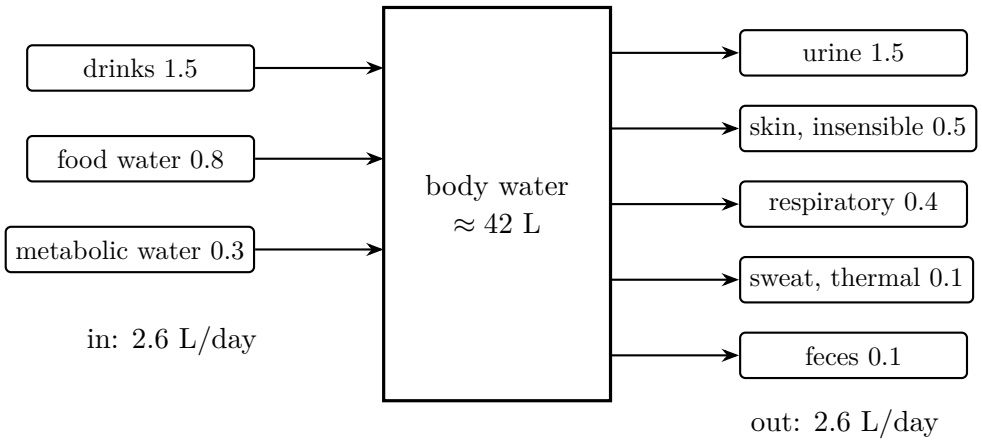


Figure 5.1: Daily water ledger of the reference adult at rest, from Table 5.1. Every stream crosses the boundary of the body, so the stores change only if the totals differ. Quantities are in liters per day.

5.4 Choosing the Boundary and Following the Dependence

A ledger for the body alone shows a body that takes in food, water, and oxygen and releases heat, carbon dioxide, urine, and sweat. A ledger for the body in a sealed suit shows a different pattern: water that left the body remains inside the boundary and has to be accounted for as a store, and the heat released at the skin has to leave the suit by a stated route. A ledger for the body, the suit, and a service facility shows that the closure of the water loop depends on inputs of energy, replacement filters, and cleaning agents entering from outside.

Worked Example 5.2 (An apparently closed loop)**.** A water recovery system returns a fraction η of the daily liquid loss L to the supply. The make-up water required per day is $(1 - \eta)L$ from Equation (5.1) applied to water alone. With $L = 2.6$ L and $\eta = 0.95$, the make-up is 0.13 L per day, and the loop seems nearly closed.

The solute ledger shows that the loop is not closed. The urine solute load of about 600 mOsm per day must leave the system or accumulate. If the system concentrates the solute to a brine of 3 $\mathrm{Osm\,kg^{-1}}$, the brine water amounts to 0.2 kg per day, and each day 0.2 kg of brine must be stored or discharged. Over a month this is six kilograms. Higher brine concentrations reduce the mass but require thermal or electrical processing that costs energy and produces crystalline residues. The energy ledger shows a further dependence: to evaporate 1.5 kg of water without heat recovery requires 1.5×2.26 MJ ≈ 3.4 MJ per day, which is about 40 percent of the metabolic energy intake of the person, so a practical system must recover latent heat by vapor compression or similar means. The loop is closed for water at the boundary of the water ledger, and open for solute, energy, and consumables at the boundary of the full ledger. Spaceflight water recovery systems operate at recovery fractions of the order of 85 to 98 percent depending on whether a brine processing stage is included [7], which is the same pattern.

5.5 Teaching the Habit of the Boundary

The habit that this chapter aims to establish has three steps. The boundary is stated first, in words and in a diagram. Every stream crossing it is then listed with its direction, its magnitude, and its dependence on conditions. Finally, the accumulation terms are examined for any quantity that grows without limit, because a quantity that grows without limit identifies the place where the apparent closure of a loop has been bought with an unaccounted external input. The remaining chapters of Part II apply this habit to separation, to control, and to scaling.

Evidence Status

Established physiology. The ledger quantities of Table 5.1 are standard physiological values [18, 5]. Demonstrated engineering. Water and air recovery systems operate on spaceflight vehicles and have documented recovery fractions [7]. Proposed design. The use of the ledger for wearable and implanted recovery systems is a framework, and the chapters that follow supply the case analyses.

Exercise 5.1. Repeat Worked Example 7.1 for the child and the large adult of Table 5.2, taking the metabolic power during exertion to scale with the same three-quarter power. State which of the two reaches a core temperature rise of 2 K first under a garment that rejects heat only by dry transfer.

Exercise 5.2. Recompute the ledger of Table 5.1 for a person engaged in sustained exertion at 400 W of metabolic power for four hours in an environment at 30 °C. State assumptions for the sweat rate, and identify which entries change by more than a factor of two.

Exercise 5.3. In Worked Example 5.2, vary the brine concentration from 2 to 6 $\mathrm{Osm\,kg^{-1}}$ and tabulate the daily brine mass and the overall water recovery fraction. Identify the concentration beyond which further gains in water recovery are smaller than the daily variation of insensible loss.

Chapter 6

Separation, Sorting, and Water Recovery

A device that handles bodily fluids must separate one stream into two or more streams of different composition. The physical principles available for doing so differ greatly in what they can select, how much energy they cost, and what condition they leave their outputs in. This chapter develops the principles needed to judge a proposed water recovery unit, with attention to the common misconception that a small centrifuge can desiccate a fluid.

6.1 Separation by Density: What a Centrifuge Can Do

A centrifuge replaces gravitational acceleration with the centripetal acceleration $a = \omega^2 r$ and acts on differences in density between a particle and its surrounding fluid. For a spherical particle of radius r_p and density ρ_p in a fluid of density ρ_f and viscosity μ, the Stokes settling velocity under acceleration a is

$$v = \frac{2}{9}\,\frac{(\rho_p - \rho_f)\, a\, r_p^2}{\mu}. \tag{6.1}$$

For a red blood cell, with $r_p \approx 3.5\ \mu\mathrm{m}$, $\rho_p - \rho_f \approx 75\ \mathrm{kg\,m^{-3}}$, and $\mu \approx 1.2\ \mathrm{mPa\,s}$, the velocity at $a = g$ is about $1.7\ \mu\mathrm{m\,s^{-1}}$, and at one thousand times gravity it is about $1.7\ \mathrm{mm\,s^{-1}}$. Cells, platelets, bacteria, precipitated crystals, and lipid droplets are removed from fluids by centrifugation on timescales of seconds to minutes.

Dissolved species behave very differently. At equilibrium in a centrifugal field, a dissolved solute of molar mass M and partial specific volume $\bar{v}$ has a concentration gradient given by the sedimentation equilibrium relation

$$\frac{\mathrm{d}\ln c}{\mathrm{d}r} = \frac{M\,(1 - \bar{v}\rho)\,\omega^2 r}{RT}. \tag{6.2}$$

The right-hand side is proportional to the acceleration and to the molar mass, and small molecules have small molar mass.

Worked Example 6.1 (Urea in a centrifuge)**.** Take urea, with $M = 0.060\ \mathrm{kg\,mol^{-1}}$ and $\bar{v}\rho \approx 0.75$ in a solution of density near $1000\ \mathrm{kg\,m^{-3}}$, in a field of $10^5 g$, so $\omega^2 r \approx 9.8 \times 10^5\ \mathrm{m\,s^{-2}}$. At $T = 310$ K,

$$\frac{\mathrm{d}\ln c}{\mathrm{d}r} \approx \frac{0.060 \times 0.25 \times 9.8 \times 10^5}{8.314 \times 310} \approx 5.7\ \mathrm{m^{-1}}.$$

Over a radial distance of one centimeter, the concentration at the outer radius exceeds that at the inner by a factor of $e^{0.057}$, or about six percent. A field of one hundred thousand times gravity, which requires an ultracentrifuge rotor, produces a six percent concentration difference for urea across a centimeter. A centrifuge small enough to be part of a body, operating at far lower acceleration, will not separate urea, sodium chloride, or glucose from water. It can sort cells and particles, and it can assist the removal of droplets and crystals from a stream. It cannot desiccate a fluid.

The example does not rule out the use of a centrifuge in a recovery unit. It specifies the role it can play. A rotating separator can remove particulates and cells from a stream before it reaches a membrane, reducing fouling, and it can separate a liquid phase from a vapor or from a gas at the exit of a phase-change stage, where the density ratio is large. In the proposals examined in this book it appears as a pre-treatment or a phase separator, and the separation of dissolved material is assigned to other mechanisms.

6.2 Separation of Dissolved Material

Membranes, phase change, and chemical affinity are the mechanisms that discriminate among dissolved species. Table 6.1 compares them on the properties that decide their use in a body-scale system.

Table 6.1: Mechanisms for separating water from dissolved material.

Mechanism	Selects on	Energy scale	Condition of output
Rotating separator	Density of particles	Low; rotation and friction	Clarified liquid; dissolved matter unchanged
Ultrafiltration	Size above about 10 nm	Low pressure, tens of kPa to a few hundred kPa	Removes macromolecules; ions and urea pass
Reverse osmosis	Solution and diffusion across a dense membrane	Exceeds the osmotic pressure, about 1.5 MPa at 600 $\mathrm{mOsm\,L^{-1}}$	Low salt; small neutral solutes such as urea pass in part
Electrodialysis	Ionic charge	Proportional to ions removed	Removes salts, leaves neutral solutes
Distillation	Volatility	Latent heat unless recovered	Very clean water; volatile organics may carry over
Ion exchange and adsorption	Specific chemistry	Consumable media	Selective, but media saturate

The thermodynamic floor for separating water from a solution is set by its osmotic pressure. For a dilute ideal solution with osmotic pressure π_0 in the feed, the minimum work per unit volume of permeate when a fraction r of the feed is recovered is

$$w_{\min} = \frac{\pi_0}{r} \ln \frac{1}{1 - r}. \tag{6.3}$$

The derivation, which integrates the osmotic pressure of the progressively concentrating feed over the permeate volume, is given in Appendix A. At 600 $\mathrm{mOsm\,L^{-1}}$ the osmotic pressure at body temperature is $\pi_0 \approx 1.55$ MPa, and for $r = 0.9$ Equation (6.3) gives about 4 $\mathrm{MJ\,m^{-3}}$, or 4 kJ per liter. Recovering 1.5 L per day requires at least 6 kJ of work, which is under one tenth of one percent of daily metabolic energy. The thermodynamic floor is not the obstacle. The obstacles are the efficiency of real pumps and membranes, which multiplies the floor by a large factor, the fouling and replacement of membranes, the incomplete rejection of urea by reverse osmosis membranes, and the handling of the brine.

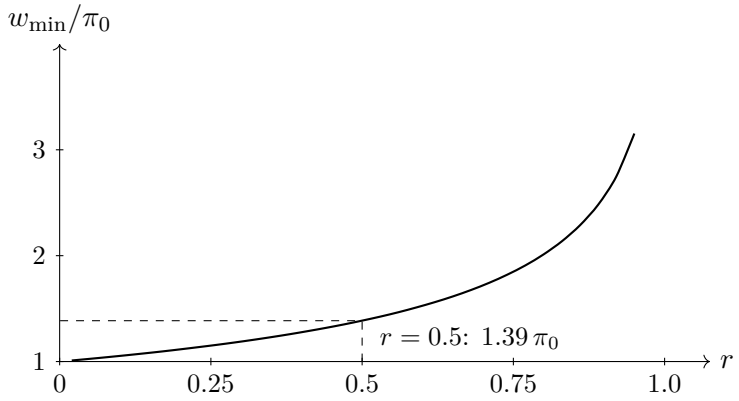


Figure 6.1: Minimum work per unit volume of recovered product in units of the feed osmotic pressure π_0, from Equation (6.3), as a function of the recovery fraction r. The work tends to π_0 as $r \to 0$ and rises without bound as $r \to 1$, which is why a recovery unit that approaches complete recovery pays a rapidly growing energy price.

6.3 Three Versions of One Unit

The central exercise of this chapter compares three placements of the same water recovery function: an internal stage within the body, a module worn in clothing, and a stationary system coupled to the body through the Umbilicus of Chapter 14. Each placement changes the constraints without changing the function.

An internal stage is limited by tissue heating. A spherical implant of radius R dissipating power at density p_d in tissue of thermal conductivity $k \approx 0.5\ \mathrm{W\,m^{-1}\,K^{-1}}$ produces a steady temperature rise at its surface of

$$\Delta T = \frac{p_d R^2}{3k}, \tag{6.4}$$

neglecting the cooling effect of blood perfusion, which helps. For an allowed rise of 1 K and a radius of one centimeter, the permitted power density is $3k\,\Delta T/R^2 = 1.5 \times 10^4\ \mathrm{W\,m^{-3}}$, or about 15 mW per cubic centimeter, so a unit of about 4 $\mathrm{cm^3}$ is limited to roughly 60 mW. Distillation with heat recovery, even at 100 $\mathrm{kJ\,kg^{-1}}$, requires about 150 kJ per day for 1.5 kg, which is 1.7 W averaged over the day. An internal thermal stage is therefore not credible at this scale, and an internal stage would have to rely on membrane processes at far lower power or on the physiological pathways of the kidney itself.

A module in clothing is free of the tissue heating limit and can reject heat to the environment, but it is limited by mass, by the area available, and by the movement of the wearer. It can carry replaceable membranes and consumable media, and it can store a brine for later removal.

A stationary system accessed through the Umbilicus can use large equipment with efficient heat recovery and can be maintained by service. It requires a connection, and its benefit is available only when the connection is made. Chapter 14 analyzes the cost of that dependence. The comparison yields the pattern that recurs through the book: the more of a function that is carried, the more constrained it is by body-scale limits on heat, volume, and maintenance, and the more of it that is external, the more the person depends on the connection.

Evidence Status

Established physiology. Stokes settling, sedimentation equilibrium, osmotic pressure, and the physics of the bioheat balance are standard [4, 6]. Demonstrated engineering. Reverse osmosis, vapor compression distillation, and rotating phase separators are used in water treatment and in spaceflight [7]. Proposed design. The three placements are design concepts developed in this book, and the internal case has been analyzed only to the level of Equation (6.4).

Exercise 6.1. Using Equation (6.3), plot $w_{\min}$ against r and find the recovery fraction at which $w_{\min}$ is twice π_0. Then modify the derivation of Appendix A for a feed whose osmotic pressure rises faster than linearly with concentration, as for a concentrated brine, and state the direction in which the minimum work changes.

Exercise 6.2. Find the rotor speed in revolutions per minute for a radius of 2 cm that gives $10^3 g$, and compute the Stokes velocity for a bacterial cell with $r_p = 0.5$ μm and a density difference of 50 $\mathrm{kg\,m^{-3}}$. State how long it would take to move the cell 1 mm.

Exercise 6.3. For each of the three placements, list the three constraints that most limit the design and one consumable that must be resupplied. Argue which placement has the shortest grace period after a failure.

Chapter 7

Physiological Control Under Changing Demand

An auxiliary organ must change what it does as the condition of the body changes, because the body's demands for water, heat rejection, and flow vary by an order of magnitude through a day. Control theory supplies the vocabulary for stating how that adjustment is achieved and when it fails. The chapter develops the minimum needed for the later case studies and works through a case in which improving water recovery in one set of conditions reduces the capacity to manage heat in another.

7.1 Feedback, Delay, and Gain

A feedback loop measures a quantity, compares it with a set point, and acts to reduce the difference. Take a first-order system with time constant τ and gain K under proportional control with gain k and a delay θ in the actuation. The characteristic equation is

$$\tau s + 1 + kK\, e^{-s\theta} = 0. \tag{7.1}$$

Stability is lost at the critical gain at which a root crosses the imaginary axis. For $\theta \ll \tau$ the crossing frequency is approximately $\omega_c \approx \pi/(2\theta)$, and the critical loop gain is

$$(kK)_{\mathrm{crit}} \approx \frac{\pi \tau}{2\theta}. \tag{7.2}$$

The result expresses a general rule: a loop with a long delay relative to the time constant of the process it controls tolerates only a low gain, and a high gain produces oscillation [1]. Physiological loops exhibit this behavior, as in oscillations of ventilation in periodic breathing and of sweating at the onset of exertion.

7.2 Interacting Loops and Disagreeing Sensors

The body controls temperature and fluid volume with several interacting loops. Core temperature is sensed in the hypothalamus and in deep tissue, skin temperature is sensed at the surface, and the thermoregulatory response depends on a weighted combination of both, with the core signal dominating at steady state. Plasma osmolality is sensed in the hypothalamus and drives thirst and the release of antidiuretic hormone, and blood volume is sensed by stretch receptors and drives sodium handling. When sensors disagree, as when the skin is hot while the core is cool, the controller must choose a weighting, and an artificial device added to the system becomes another sensor or actuator that the weighting did not anticipate.

Two consequences follow for auxiliary organs. A device that acts on a variable already regulated by the body, for example by recovering water, introduces a second controller of that variable, and the two may conflict. A device that changes a path of heat or water loss changes the plant that the native controller regulates, so the controller's gain and delay are no longer matched to it.

7.3 A Case of Competing Objectives

Worked Example 7.1 (Water recovery against heat rejection)**.** A person of mass 70 kg and skin area $A = 1.8\ \mathrm{m}^2$ works at a metabolic rate of 500 W and an external mechanical output of 100 W, so the heat to be rejected is 400 W. Take a combined dry heat transfer coefficient $h = 8\ \mathrm{W\,m^{-2}\,K^{-1}}$, so $hA = 14.4\ \mathrm{W\,K^{-1}}$, a mean skin temperature of 34 °C, and a latent heat of vaporization of sweat of $2.43\ \mathrm{MJ\,kg^{-1}}$, so that a sweat rate of $1\ \mathrm{L\,h^{-1}}$ that is fully evaporated removes about 675 W.

Open skin. At an air temperature of 20 °C the dry loss is $14.4 \times 14 \approx 202$ W, and the evaporation needed to balance the load is 198 W, or $0.29\ \mathrm{L\,h^{-1}}$. At 35 °C the dry term is a gain of 14 W and the required evaporation is 414 W, or $0.61\ \mathrm{L\,h^{-1}}$.

Closed garment. Now let a garment hold all the sweat vapor so that nothing leaves the garment but heat. The boundary is the outer surface of the garment, of the same area and with the same coefficient, so $hA = 14.4\ \mathrm{W\,K^{-1}}$ for that surface only, and the surface is credited once. In steady state the full 400 W must leave through it as dry heat, which requires a surface-to-air difference of $400/14.4 \approx 27.8$ K. At an air temperature of 20 °C the surface would have to reach 47.8 °C. Limiting it to a touch-safe 45 °C gives a rejection of $14.4 \times 25 = 360$ W and a deficit of 40 W. At 35 °C the same surface rejects only $14.4 \times 10 = 144$ W and the deficit is 256 W. A closed garment of this area therefore balances 400 W only for air temperatures below about 17 °C, since $45 - 27.8 = 17.2$.

Consequence under a lumped model. Treat the body as a single heat capacity of $70 \times 3470 \approx 2.4 \times 10^5\ \mathrm{J\,K^{-1}}$ and let the deficit be stored. A rise of 2 K, about 4.9×10^5 J, takes $4.9 \times 10^5/40 \approx 12{,}000$ s, or 3.4 hours, at 20 °C, and $4.9 \times 10^5/256 \approx 1900$ s, about half an hour, at 35 °C. These are times for a lumped temperature to rise under the stated model. They are not predictions of core temperature or of safe working time, which would require a model with distinct core and skin compartments and with the control of Equation (7.1).

Partial recovery. If the garment lets vapor escape, the exported latent heat is $400 - 14.4\,(T_s - T_a)$, and the water that can be retained is the sweat produced minus the water exported at that rate. At 20 °C with a surface at 45 °C the export needed is 40 W, or $0.06\ \mathrm{L\,h^{-1}}$, so most of the sweat, about four fifths of the $0.29\ \mathrm{L\,h^{-1}}$ of the open case, can be retained. At 35 °C the export needed is 256 W, or $0.38\ \mathrm{L\,h^{-1}}$, which is more than half of the $0.61\ \mathrm{L\,h^{-1}}$ the body would sweat. The same surface serves two functions, water retention and heat rejection, and has a single budget to rejection. The recovery fraction is therefore a function of the environment and cannot be specified as a constant.

The Running Problem Returns

For the running problem of Chapter 4 the heat to be rejected during the two hours is $300 - 75 = 225$ W. In e_1 the closed garment G with a surface at 45 °C rejects 360 W, a surplus of 135 W, so the heat balance closes with all water retained. In e_2 it rejects 144 W and the deficit is 81 W, which raises the lumped temperature by 2 K in $4.9 \times 10^5 / 81 \approx 6000$ s, or 100 minutes, so under this model the garment leaves the viable set before the two hours are complete unless it opens an evaporative path. Water and heat are now checked together, and the answer for G depends on the environment.

7.4 Design Consequences

A water recovery device that sits in the path of heat rejection must be able to be switched out of that path when the environment demands it. The control problem is to detect from available signals, which include skin and core temperature, ambient temperature, humidity, and the rate of heat storage, that the dry rejection is inadequate, and to open an evaporative path in time. The detection delay and the rate at which the garment can change its resistance determine whether the switch is made before the core temperature has risen past the margin. A failure to include the switch is an example of a specification failure of the type described in Chapter 4, because a design that satisfies the water ledger fails the specification when the environment moves beyond the range for which it was tested.

Evidence Status

Established physiology. Thermoregulation, the latent heat of sweat, the energy balance of the body, and the principles of delay-limited feedback are standard [18, 16, 1]. Demonstrated engineering. Liquid-cooled and vapor-managing garments exist and have been used in protective and aerospace work [7]. Proposed design. The sweat-recovering garment of Worked Example 7.1 is a design concept, and the numbers are illustrative parameters that fix the order of magnitude.

Exercise 7.1. For the closed garment of Worked Example 7.1, find the outer surface area that balances 400 W at an air temperature of 35 °C with a surface at 45 °C, and discuss what limits the area that a garment can have.

Exercise 7.2. The body has a thermal time constant of roughly $\tau = C/(hA)$ when sweating is absent. Evaluate it for the parameters above. Use Equation (7.2) with a delay of five minutes to find the maximum loop gain of a proportional sweat controller.

Exercise 7.3. Propose a decision rule for opening the evaporative path of a recovering garment, stating the signals used and the threshold, and identify one false alarm and one missed detection that the rule would produce.

Chapter 8

Scaling, Redundancy, and Repair

The scale at which a function is performed decides how much of it can be done, how reliably, and with what waste heat. This chapter derives the geometric laws that govern shrinking devices, compares a centralized organ with a population of small units, and develops the threshold model of the hepastitium, in which the whole organ works if enough of its nodes work.

8.1 Geometric Scaling

For a device of linear dimension L, surface area scales as L^2 and volume as L^3. Quantities that depend on the volume, such as the mass of stored reserve, grow faster than those that depend on the surface, such as heat rejection, which is the basis of the observation that large bodies cool slowly and small bodies cool quickly. For a cube dissipating power density p_d in an environment with heat transfer coefficient h, the temperature rise is

$$\Delta T = \frac{p_d L^3}{6 h L^2} = \frac{p_d L}{6h}, \tag{8.1}$$

so for a fixed power density, reducing the size lowers the temperature rise. For a spherical implant in conducting tissue, Equation (6.4) gives $\Delta T \propto p_d R^2$, and the effect is stronger. Shrinking a component at constant power density helps its thermal state, whereas shrinking it while raising its power density, as when a miniaturized pump must deliver the same flow, can worsen it.

Metabolic scaling adds a second law. Across mammals the basal metabolic rate scales approximately as body mass to the three-quarter power [21, 34], and a stored reserve of fixed fractional mass scales as the first power. Taking mass proportional to L^3, a reserve $S \propto L^3$ is drawn down at a rate $q \propto L^{9/4}$, so the endurance time $t_e = S/q$ scales as

$$t_e \propto L^{3/4}. \tag{8.2}$$

Halving the linear dimension of an animal with a fixed reserve fraction therefore reduces its endurance without resupply to about $0.5^{0.75} \approx 0.6$ of the original. The relation is an interspecies regularity and not a law for any redesigned human, but it is the baseline against which claims about smaller bodies must be compared in Part V.

The Running Problem Returns

Architecture H of the running problem recovers fluid by a mesh of nodes. Suppose its function is adequate when at least 60 of 100 nodes work, each surviving the day with probability 0.7 independently, and suppose a shared power and coordination layer fails with probability $q = 0.01$ in the same day. Then the endpoint reliability is $0.9875 \times 0.99 \approx 0.978$,

so the shared layer already costs about one percent of reliability. In environment e_3, with all external support withdrawn for three days, H is exposed to three days of node failure with no replacement, so the endpoint reliability must be recomputed for the longer interval and the first-crossing event matters. The garment G has no node population and fails differently, and U has no grace period beyond what the body retains, as Chapter 14 shows.

8.2 Centralized and Distributed Units

A centralized unit is one device of capacity C. A distributed unit is a population of n small devices of individual capacity C/n. The centralized unit has lower overhead per unit capacity and fewer interfaces. The distributed unit tolerates the loss of members and can place support near the tissue that needs it. A population also has a different failure structure, and the right model is one that states how failures are related.

8.3 The Hepastitial Threshold Model

Suppose that the organ consists of n nodes, that each node is working at the end of a stated horizon with probability r, and that the organ's function is adequate when at least k nodes are working. If nodes fail independently, the probability that the function is adequate is

$$R_{\mathrm{ind}}(n,k,r) = \sum_{j=k}^{n} \binom{n}{j} r^{j} (1-r)^{n-j}. \tag{8.3}$$

A single node with $r = 0.9$ gives $R = 0.9$. A population of $n = 100$ nodes requiring $k = 60$ with individual $r = 0.7$ gives $R_{\mathrm{ind}} \approx 0.9875$, so a population of individually less reliable parts can exceed the reliability of a single better part.

Proposition 8.1 (A marginal node)**.** Fix k and let R_n be the reliability of Equation (8.3) for n independent nodes. Then $R_{n+1} - R_n = r\,\mathbb{P}(X_n = k-1) \geq 0$, where X_n is the number of working nodes among the first n. Under the common-layer model of Equation (8.4), the corresponding gain is $(1-q)\,r\,\mathbb{P}(X_n = k-1)$.

Proof. Couple the two systems by giving the $(n+1)$-node system the outcomes of the first n nodes and one further independent node. The added node changes the outcome only when the first n nodes have exactly $k-1$ working, so that the new node, working with probability r, brings the count to k. Hence the increase in the probability of adequacy is $r\mathbb{P}(X_n = k-1)$. The common layer multiplies every outcome by the same factor $1-q$. □

The reliability calculated here is the probability that the threshold is met at the stated time. With repair or replacement a system can fall below the threshold in the interior of the interval and recover by the endpoint, and continuous support calls for a different event, the probability that the number of working nodes $N(t)$ stays at or above k throughout, $\mathbb{P}(\inf_{0\leq t\leq T} N(t) \geq k)$. The endpoint calculation is the introductory result, and modeling the first crossing of the threshold is the subject of an exercise.

Worked Example 8.1 (Relaxing independence)**.** The result of Equation (8.3) depends on independence, which rarely holds. Three relaxations are instructive.

First, a shared power or coordination layer that fails with probability q and disables all nodes gives

$$R_{\mathrm{sys}} = (1 - q)\, R_{\mathrm{ind}}(n, k, r), \tag{8.4}$$

so the reliability cannot exceed $1 - q$ however large n is. Adding nodes still helps, as Proposition 8.1 shows, but every gain is multiplied by $1 - q$ and the limit is $1 - q$, so the design problem moves to the reliability of the common layer.

Second, nodes may share a manufacturing or environmental quality, so that r is itself a random variable. With r drawn from a Beta distribution with mean 0.7, the system reliability for $n = 100$ and $k = 60$ is about 0.77 for a broad distribution (Beta with parameters 7 and 3) and about 0.94 for a narrow one (parameters 70 and 30), compared with 0.9875 under independence. Heterogeneity of quality across a population that is correlated within each device lot is therefore as important as the mean.

Third, maintenance can restore nodes. If failed nodes are replaced at a rate that exceeds the failure rate, the number of working nodes settles near a steady value and the question becomes whether the maintenance process itself is available, which links the model to the question of dependence in Chapter 20.

8.4 Coverage and Connectivity

Three different quantities have to be kept apart: the expected uncovered fraction of the tissue, the probability that every region is covered, and the probability that the working nodes form a connected network. The Poisson relation below gives only the first, for a homogeneous field of nodes in unbounded space. It does not give the other two, and for a finite tissue region it ignores edge effects, which reduce coverage near the boundary. A finite organ model needs an explicit event, such as coverage of every critical region together with a path from each to a working relay.

Two geometric conditions must hold for a population of nodes to act as an organ. Every region of tissue to be supported must lie within reach of a working node, and the working nodes must form a connected network if they depend on one another for coordination. If nodes are scattered as a Poisson process of density λ and each supports a sphere of radius r_c, the fraction of tissue left unsupported is

$$f_{\mathrm{unc}} = \exp\!\left(-\tfrac{4}{3}\pi r_c^3 \lambda\right). \tag{8.5}$$

Connectivity is a different condition. For nodes that communicate within a range r_{com}, a three-dimensional random geometric network in an infinite medium percolates, in the sense of containing an infinite connected component, only when the mean number of neighbors per node exceeds a critical value near 2.7 [31]. This is an approximation to be investigated in finite geometry, where a spanning cluster of a finite organ is a different event. Coverage of every region is therefore compatible with a mesh that has fallen into disconnected islands, and the two conditions need to be checked separately. As nodes fail, λ decreases, and the order in which coverage and connectivity are lost determines whether the organ degrades gradually or in a step.

Evidence Status

Established physiology. Geometric scaling, the three-quarter-power law for metabolic rate across species, and the k-out-of-n reliability formulas are established [21, 34, 27, 26]. The percolation threshold is a result of the theory of random media [31]. Demonstrated engineering. Redundant populations of small devices are used in other engineering fields, including sensor networks and distributed storage. Proposed design. The hepastitial threshold model is a design analysis, and its parameters for an implanted population have not been measured.

Exercise 8.1. Compute R_{ind} for $n = 20$, $k = 12$ and $r = 0.7$, and for a single device with the reliability needed to match it. Discuss what the comparison implies about the use of many small parts.

Exercise 8.2. Using Equation (8.4), find the largest q that permits a system reliability of 0.95 when $R_{\mathrm{ind}} = 0.9875$ is the independent term. Interpret the answer in terms of the design of the power and coordination layer.

Exercise 8.3. A failed node is replaced after a random delay. Write the event that the number of working nodes stays at or above k throughout an interval $[0, T]$, and explain why its probability is below that of the endpoint event of Equation (8.3). Propose a simulation to estimate it.

Exercise 8.4. For a tissue region of 1 cm^3 supported by nodes with $r_c = 2$ mm, find the density λ at which the unsupported fraction falls below one percent. Estimate the number of nodes needed to cover a liter of tissue.

Part III

Reassigning the Work of the Body

Chapter 9

Circulation, Respiration, and the Moving Body

Circulation and respiration deliver oxygen to tissue and remove carbon dioxide, and they do so under demands that change by a factor of ten or more between sleep and sustained exertion. The question for an auxiliary organ is where in the system pumping and gas exchange can be added or relocated, which functions benefit from support close to the tissue, and how the same system behaves at rest, in exertion, and when its supply is interrupted. The answer depends on location as much as on nominal capacity.

9.1 Oxygen Delivery and Extraction

The oxygen content of arterial blood is

$$C_{aO_2} \approx 1.34\,[\mathrm{Hb}]\,S_{aO_2} + 0.003\,P_{aO_2}, \tag{9.1}$$

in milliliters of oxygen per deciliter when the hemoglobin concentration is given in grams per deciliter and the oxygen partial pressure in millimeters of mercury, and the dissolved term is small. Oxygen delivery is the product of cardiac output and arterial content, and consumption is the fraction of delivery that tissue extracts.

Worked Example 9.1 (Delivery at rest and in exertion)**.** For $[\mathrm{Hb}] = 15\ \mathrm{g\,dL^{-1}}$ and $S_{aO_2} = 0.98$, the content is about $19.7\ \mathrm{mL\,dL^{-1}}$, or $197\ \mathrm{mL\,L^{-1}}$. At rest, with a cardiac output of $5\ \mathrm{L\,min^{-1}}$, delivery is about $985\ \mathrm{mL\,min^{-1}}$. Consumption of $250\ \mathrm{mL\,min^{-1}}$ then corresponds to an extraction of about 25 percent. In heavy exertion, with a consumption of $3000\ \mathrm{mL\,min^{-1}}$ and a cardiac output of $20\ \mathrm{L\,min^{-1}}$, delivery is about $3940\ \mathrm{mL\,min^{-1}}$ and the extraction is $3000/3940 \approx 76$ percent. Two quantities rose, flow by a factor of four and extraction by a factor of three, to produce a twelvefold increase in consumption. A device that supplied only flow, and left extraction unchanged at 25 percent, would need an output of $3000/(0.25 \times 197) \approx 61\ \mathrm{L\,min^{-1}}$ to match the same consumption, which is far beyond the physical capacity of any wearable pump.

The example shows that circulatory support cannot be assessed from flow alone. The body uses local mechanisms, including increased capillary density in active muscle, a rightward shift of the hemoglobin dissociation curve in warm acidic tissue, and the opening of vascular beds, to raise extraction where it is needed. Any added pump must respect these mechanisms or be designed to supply them.

9.2 Local Support and the Limits of Diffusion

Oxygen moves from capillary to cell by diffusion over distances of order 100 μm. With a tissue diffusivity of about 2×10^{-9} m^2 s^{-1}, the characteristic time for diffusion over a distance x is $t \approx x^2/(2D)$, which gives 2.5 s for $x = 100$ μm and 250 s for $x = 1$ mm. This explains why every tissue needs a capillary within a fraction of a millimeter and why a supplementary device that serves a region of tissue must either match that density or operate by carrying the supply to the capillaries that already exist. Local support is therefore most effective when it acts on convective flow, as muscle pumps and assist devices do, or on the gas content of blood, as oxygenators do, and not when it tries to supply cells directly across tissue distances greater than a fraction of a millimeter.

Location matters in a second sense. The calf muscle pump returns venous blood against gravity and works only when the person moves, so a function that is supported by movement is lost when movement stops. A distributed assist that sits near the limb can restore this function at rest, whereas a central pump cannot reproduce the local pressure pattern. A function can have the same nominal capacity in two places and have a different effect.

9.3 How Close Must Support Be? The Krogh Cylinder

The diffusion time estimate above can be sharpened. In the Krogh model each capillary of radius R_c supplies a coaxial cylinder of tissue of radius R_t that consumes oxygen at a uniform volumetric rate q. With no flux across the outer surface of the cylinder, steady diffusion with coefficient D gives the largest drop in oxygen concentration, at the outer radius, as

$$\Delta C = \frac{q}{4D}\left[R_t^2 \ln\frac{R_t}{R_c} - \frac{R_t^2 - R_c^2}{2}\right]. \tag{9.2}$$

Tissue is adequately supplied when ΔC is less than the concentration available at the capillary end.

Worked Example 9.2 (Spacing in exercising muscle)**.** Take $D = 2 \times 10^{-9}$ m^2 s^{-1}, $R_c = 3$ μm, and an oxygen concentration of about 0.054 mol m^{-3} at the capillary, the amount dissolved at a partial pressure near 40 mmHg. Resting muscle consumes about 4×10^{-3} mol m^{-3} s^{-1}, which at $R_t = 40$ μm gives $\Delta C \approx 1.5 \times 10^{-3}$ mol m^{-3}, far inside the supply. Heavily exercising muscle consumes about 0.15 mol m^{-3} s^{-1}, and Equation (9.2) gives $\Delta C = 0.011$ at $R_t = 20$ μm, 0.031 at 30 μm, and 0.063 at 40 μm, which exceeds the supply. The limiting radius is about 38 μm. In exertion, the cells must lie within a few tens of micrometers of a capillary. Hemoglobin-like carriers in the tissue facilitate transport and relax the bound, and the number uses only dissolved oxygen, so it is conservative. A supplementary node that is meant to support tissue by delivering oxygen or a metabolite directly must therefore sit within tens of micrometers of the cells it serves, or must act through the existing vasculature. The coverage radius $r_c = 2$ mm assumed in Chapter 13 is justified only for functions that do not depend on diffusion from the node.

9.4 Interruption

The body stores about 1.5 L of oxygen in blood, myoglobin, lungs, and tissue water. At a resting consumption of 0.25 $\mathrm{L\,min^{-1}}$ this supports about six minutes, and at 3 $\mathrm{L\,min^{-1}}$ about thirty seconds, if consumption is not reduced. The brain consumes roughly a fifth of resting oxygen and tolerates interruption for only a few minutes before injury. For circulatory and respiratory support, the grace period τ_g defined in Chapter 4 is therefore measured in minutes, and a design that depends on an external connection for these functions has no practical fallback unless the connection is redundant or the device is carried with an independent source. This contrasts with water and food, for which the grace period is days to weeks, and it explains the ordering of functions by dependence in the later chapters: gas exchange and perfusion are carried, and water and nutrition may be accessed.

Evidence Status

Established physiology. The oxygen content relation, typical values for flow and extraction, diffusion distances, and oxygen stores are standard physiology [18, 5]. Demonstrated engineering. Ventricular assist devices and extracorporeal oxygenators are in clinical use, but wearable lung assist is limited to investigational prototypes [22]. Proposed design. Distributed local assist, as discussed in Chapter 13, is a proposal.

Exercise 9.1. Repeat Worked Example 9.1 for a person at altitude with $S_{aO_2} = 0.85$ and the same hemoglobin concentration. Find the cardiac output required for a consumption of 3 $\mathrm{L\,min^{-1}}$ at an extraction of 76 percent.

Exercise 9.2. A wearable pump is designed to add 2 $\mathrm{L\,min^{-1}}$ to the circulation of a person with impaired cardiac output. Identify the quantities that must be measured to decide whether this addition permits sustained exertion, and the assumption about extraction on which the decision depends.

Chapter 10

Digestion, Renal Function, and Metabolic Exchange

The body's handling of food, water, and waste appears to be one system, but it consists of four separable functions. The first is the intake of food and drink. The second is the making of nutrients available in a form that can cross the gut wall. The third is the regulation of water, electrolytes, and acid, which the kidney does continuously. The fourth is the handling of wastes, which are removed in urine, feces, breath, and sweat. An auxiliary organ can be proposed against any of these, and the cost of a proposal differs greatly according to whether it supplements one operation or takes on the regulation that surrounds it.

10.1 Regulation Is More Than Removal

The kidney is not a filter whose output is fixed by the composition of blood. It regulates the composition of blood by changing its output. Urine osmolality ranges from about 50 $\mathrm{mOsm\,kg^{-1}}$ in water diuresis to about 1200 $\mathrm{mOsm\,kg^{-1}}$ at maximal concentration. The minimum urine volume is set by the solute that must be excreted and the highest concentration that can be reached. For a solute load of 600 mOsm per day,

$$V_{\min} = \frac{600\ \mathrm{mOsm}}{1200\ \mathrm{mOsm\,kg^{-1}}} = 0.5\ \mathrm{kg\ per\ day}. \tag{10.1}$$

This obligatory loss is the water cost of excreting solute with the native kidney and sets the lower bound on water turnover in a person who is not recovering water elsewhere. Recovery of water from urine by an external device changes the cost: the carrier water is returned, and the obligatory loss is replaced by the cost of treating a brine.

10.2 Supplementing an Operation versus Assuming the Regulation

Hemodialysis supplements the removal of urea, potassium, and excess fluid on an intermittent schedule. Its action is open loop within a session, and the patient's blood chemistry follows a sawtooth pattern between sessions, with solute accumulating on the days without treatment and being removed rapidly during treatment. Continuous or wearable approaches reduce the sawtooth by operating over a longer period at a lower rate, and a wearable artificial kidney has been tested in a small clinical trial [17]. None of these assumes the regulatory function, that is, the adjustment of excretion to intake, which must be handled by the patient through diet and fluid restriction or by the device's own control law.

Nutrient delivery by a route that bypasses the gut presents a parallel case. Parenteral nutrition delivers the products of digestion directly to the blood. It supplements the

operation of making nutrients available, and it removes the stimulus that maintains the gut mucosa, whose cells depend on luminal contact with food and on local hormonal signals. A design that delivers nutrients through an interface such as the Umbilicus of Chapter 14 has to say what replaces those signals, and whether the gut and its resident organisms are maintained for later resumption of ordinary feeding.

Worked Example 10.1 (Returning the carrier)**.** Let a person excrete a solute load $S = 600$ mOsm per day. Under the native kidney the urine volume is at least $V_{\min} = 0.5$ kg and typically 1.5 kg at ordinary hydration. Suppose an external unit receives the urine and returns purified water with recovery fraction η while concentrating the solute to a brine at $c_b = 3$ Osm kg^{-1}. The brine mass is $S/c_b = 0.2$ kg per day. The net water loss from the system is that brine water plus the unrecovered fraction of the remaining carrier. For 1.5 kg of urine, the maximum recovery fraction is $(1.5 - 0.2)/1.5 \approx 0.87$ when the brine is the only loss, and any lower η adds a further loss of $(0.87 - \eta) \times 1.5$ kg.

If the person also reduces the carrier volume by concentrating the urine natively to $V = 0.6$ kg, the maximum recovery fraction falls to $(0.6 - 0.2)/0.6 \approx 0.67$, though the absolute water recovered is less, 0.4 kg per day. The recovered water is useful in either case, but the burden on the device depends on how much carrier the kidney supplies. The recovery unit and the kidney are therefore coupled through the carrier volume, and a design that treats the kidney as a fixed source mispredicts the demand.

10.3 Intermittent Clearance as a Periodic Problem

The sawtooth of intermittent support can be computed. Let a solute be generated at rate G in a single compartment of volume V, and cleared at rate K during sessions of duration t_s that repeat with period T. During a session the concentration obeys $V\dot{C} = G - KC$, and between sessions $V\dot{C} = G$. Writing $a = Kt_s/V$ and $g = G/V$, the periodic solution has a pre-session concentration

$$C_{\mathrm{pre}} = \frac{G}{K} + \frac{g\,(T - t_s)}{1 - e^{-a}}, \tag{10.2}$$

and a post-session concentration $C_{\mathrm{post}} = G/K + (C_{\mathrm{pre}} - G/K)e^{-a}$. The first term is the concentration that continuous clearance at rate K would hold.

Worked Example 10.2 (Peak and mean for urea)**.** Take $G = 0.35$ mol per day, $V = 40$ L, $K = 200$ mL min^{-1} for $t_s = 4$ h, and a gap with $T = 48$ h. Then $a = 1.2$ and $g = 6.1\ \mu$mol L^{-1} min^{-1}, and Equation (10.2) gives $C_{\mathrm{pre}} \approx 24$ mmol L^{-1} and $C_{\mathrm{post}} \approx 8$ mmol L^{-1}, the first term being 1.2 mmol L^{-1}. A numerical integration agrees, and gives a time-average over the cycle of about 16 mmol L^{-1}. A continuous device would hold this same time-average with a clearance of $G/16$ mmol L$^{-1} \approx 15$ mL min^{-1}, which is one thirteenth of the clearance used by the intermittent schedule, and it would remove the peak of 24. Continuous operation matches the mean with far lower throughput, which is the engineering case for wearable devices, and the price is that it has to be worn, powered, and maintained continuously, which is the dependence analyzed in Chapter 20.

10.4 Food, Heat, and the Timing of Energy

The energy released by food is a slow source against fast loads. The body stores energy as glycogen, about 1.5 MJ in muscle and liver, and as fat, of the order of hundreds of megajoules, and the supply to tissues is controlled by hormones that act over minutes to hours. An artificial pancreas that closes the loop between glucose sensing and insulin delivery is an example of supplementation of regulation that has both sensing and actuation externalized, and it illustrates the difficulty: the sensor reads interstitial glucose, which lags blood glucose by several minutes, and the insulin acts with a delay of tens of minutes. The loop is delay-limited in the sense of Equation (7.2), and its design is an exercise in conservative control.

Evidence Status

Established physiology. The range of urine concentration, solute loads, and the regulation of excretion are standard [5, 18]. Demonstrated engineering. Hemodialysis and closed-loop insulin delivery are in clinical use, and wearable dialysis devices have been tested in small trials [17]. Proposed design. Water recovery by external or worn equipment from urine and the coupling of such equipment to the native kidney, as developed here, are design analyses.

Exercise 10.1. Using Equation (10.2), find the session length at which C_{pre} falls to 20 $\mathrm{mmol\,L^{-1}}$ for the parameters of Worked Example 10.2, and the clearance of a continuous device that gives the same peak.

Exercise 10.2. Solve Worked Example 10.1 for brine concentrations from 2 to 6 $\mathrm{Osm\,kg^{-1}}$ and a native urine volume of 1.0 kg. Tabulate the maximum recovery fraction.

Exercise 10.3. For a person on intermittent dialysis three times a week, sketch the time course of plasma potassium over a week. Describe the control strategy of a continuous wearable device that would remove the sawtooth, and identify the sensor on which it would depend.

Chapter 11

Structural Organs and Engineered Integument

Structural organs bear load, resist impact, transmit force, and maintain shape. Teeth, bone, cartilage, tendon, and skin are made of the same few materials arranged for different tasks, and engineered substitutes or additions can be placed at different locations to do the same tasks. The questions in this chapter are how load is carried at different dimensions, what restricts the rate at which a structure can be remodeled, and what an external layer can do that an internal one cannot.

11.1 Rigidity, Load, and Scale

For a beam of length ℓ under a bending moment M, with section modulus Z, the peak stress is $\sigma = M/Z$. Under geometric similarity, in which all dimensions of a body are scaled by λ, the load W scales as λ^3, the moment as $W\ell \propto \lambda^4$, and the section modulus as λ^3, so

$$\sigma \propto \lambda. \tag{11.1}$$

A larger body made of the same materials in the same proportions loads its bones more heavily, and a smaller one more lightly. This is the geometric basis of the observation that structure must become relatively thicker as animals grow larger [27]. For design, it means that a person who is made smaller by uniform scaling has structural margin to spare, and that a person made larger needs a disproportionate increase in cross-section.

11.2 Remodeling and Its Rate Limits

Bone adapts to load. The mechanostat hypothesis states that bone mass is increased when strain exceeds a threshold and reduced when strain falls below a lower one [15]. Distraction osteogenesis demonstrates that bone can be lengthened in a controlled way: when the two segments of an osteotomized bone are separated at about one millimeter per day, divided into several small increments, new bone forms in the gap [20]. Lengthening at this rate is limited by the adaptation of the soft tissues around the bone, which include muscle, tendon, nerve, and vessel, and not by the bone itself.

Worked Example 11.1 (Continuous skeletal adjustment)**.** The proposal of continuous skeletal adjustment replaces episodic lengthening with a slow, continuous change of segment length under feedback. Consider a change of 10 cm in a limb segment. Under distraction at 1 mm day^{-1} the change takes about 100 days, and under a continuous adjustment of 1 cm per year it takes ten years, at a rate of 0.027 mm day^{-1}.

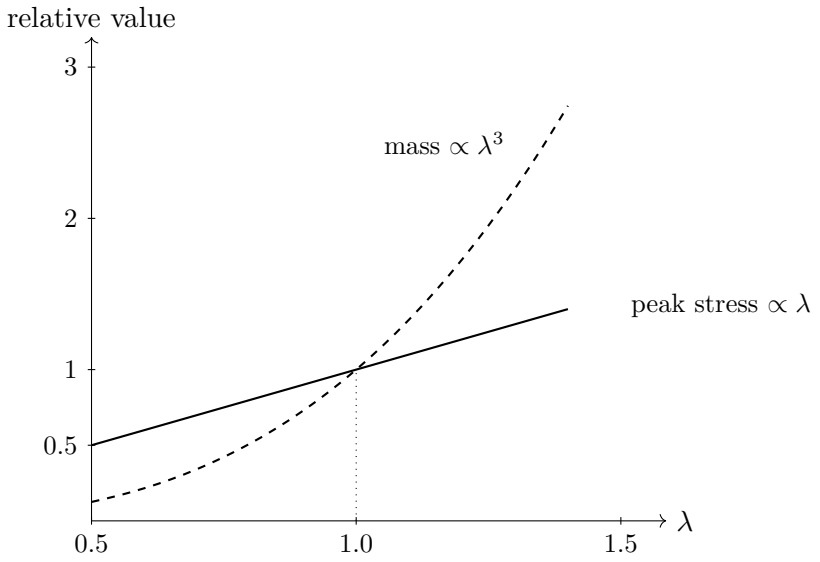


Figure 11.1: Geometric scaling of a body under similarity, from Equation (11.1). The curves are normalized to 1 at $\lambda = 1$. Mass rises as the cube of the scale and the peak stress in a bone rises only linearly, but the stress is nonetheless larger in a larger body made of the same materials.

Muscle fibers must lengthen in step. With a resting sarcomere length of about 2.5 μm, a lengthening of 1 mm per day requires the addition of 1 mm/2.5 μm = 400 sarcomeres in series per fiber per day, whereas 0.027 mm per day requires about eleven. The slower rate lies within the range of what muscle adds during ordinary growth, which suggests that the adjustment could be absorbed by remodeling and not by stretch. Nerves and vessels are less able to lengthen, and their tolerance to strain, of the order of ten percent before conduction or flow is impaired, defines a bound on the rate that depends on the starting length. A segment of 40 cm with a nerve of the same length accommodates about 4 cm of cumulative strain only if remodeling of the nerve occurs at the same time.

The movements that the person can still perform also change. A longer segment changes the moment arms of muscles, the range of joint motion, and the coordination of gait, and the sensorimotor system must relearn the mapping between command and movement, as taken up in Chapter 19. The specification of Chapter 4 thus has to include the preservation of movements, with the changing geometry as a disturbance in $\mathcal{D}$.

11.3 Skin as an Organ and the External Alternative

The skin is the largest organ, with an area of about 1.8 m^2 and a mass of several kilograms. It is a mechanical organ with a nonlinear stress-strain curve that stiffens at large strain [29], a thermal organ that varies its blood flow and sweating, a sensory organ with a dense array of receptors, and a barrier whose outer layer is renewed over about a month. These functions are mingled in one tissue. An engineered integument is an attempt to unbundle them. A layer worn outside the skin can carry mechanical protection, thermal and fluid handling, and electronics without being part of the body, and it can be replaced entire. A layer implanted beneath the skin can carry sensing and local support but must be tolerated

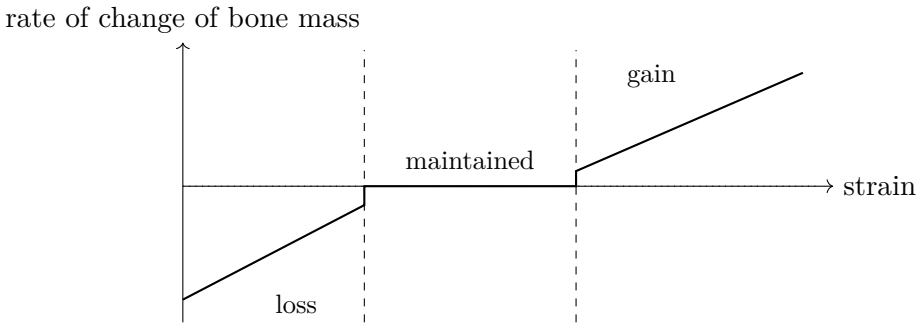


Figure 11.2: A schematic of the mechanostat hypothesis [15]: bone mass falls when strain is below a lower threshold, is maintained between the thresholds, and rises above an upper threshold. The curve shows the qualitative shape and carries no numerical values.

for years by living tissue.

The external placement has a distinctive advantage in repair, which is that the layer can be removed and renewed without surgery, and a distinctive cost, which is that it must be maintained because the body will not do it. The permanent integument introduced in Chapter 15 is a limit case of this trade-off, in which routine renewal is built into the design and the person's own skin becomes the boundary on which the layer rests.

Evidence Status

Established physiology. Scaling of stress, the mechanostat, and the viscoelastic properties of skin are established [15, 29, 27]. Demonstrated engineering. Distraction osteogenesis is a clinical procedure with documented rates and limits [20], and protective garments with impact-absorbing layers are in use. Proposed design. Continuous skeletal adjustment, under a closed loop and with accommodation of soft tissue by remodeling, is a design proposal [9].

Exercise 11.1. For a uniformly scaled person, find the factor by which the peak bone stress changes if all linear dimensions are reduced by 20 percent, and the factor by which the mass changes. Discuss why the stress change does not translate directly to a change in fracture risk.

Exercise 11.2. List the movements of a person that depend on limb segment length, and for each propose a measurement that would detect loss of function during a continuous adjustment of 1 cm per year.

Chapter 12

Sensation, Actuation, and the Fractal Brain Keel

Sensation, voluntary control, and the anatomical site of an exchange are three descriptions of an interface that are easily confused. A sensor can be moved without changing the sensation it produces, and a motor command can be delivered to an actuator at a different site from the muscle it formerly drove. This chapter separates the three descriptions, and then turns to the Fractal Brain Keel, a proposal about the geometry available for cortical growth, which is evaluated on separate criteria.

12.1 Three Separable Descriptions

The first description is the transduction of a physical quantity into a neural signal, which occurs at a receptor at a particular location. The second is the pattern of neural activity that the brain interprets as a sensation or generates as a command, which depends on the pathways and the maps in cortex. The third is the site at which a material or energy exchange with the environment occurs. These can be rearranged. Tactile-visual substitution systems deliver an image as a pattern of touch on the skin, and users of such systems learn to localize objects in space [3]. Amputees may experience sensation in a missing limb when the territory of a neighboring cortical region expands into the territory of the missing one [8]. These findings show that the relation between site, signal, and sensation is learned and adaptable. A redesign that moves a sensor, for example from the hand to the integument at the waist, must therefore plan the training that remaps the sensation and the extent to which the person’s own body schema accepts the change.

Worked Example 12.1 (Relocating a tactile sensor)**.** The number of distinguishable sites on a patch of skin is roughly its area divided by the square of the two-point discrimination threshold [33]. At the fingertip the threshold is about 2 mm, so a pad of 1 cm^2 resolves about 25 sites. On the back the threshold is of the order of 4 cm, so the whole torso, taken as 0.5 m^2, resolves about 310. The density of resolvable sites is thus about 25 per square centimeter at the fingertip and 0.06 per square centimeter on the torso, a ratio of several hundred. Moving a sensor from the hand to the waist gains area and loses resolution, and the information that can be delivered to the person is set by the number of resolvable sites and the rate at which each can be updated. A relocated sensor is therefore a different channel, with its own capacity, and training can improve the use of the channel without increasing the number of sites the skin can distinguish.

12.2 Cortical Surface Area and Its Constraints

The human cerebral cortex has an area of about 2000 cm^2 and a thickness of about 2.5 mm, which gives a volume of gray matter of about 500 cm^3. It is folded, so only a part of the area is exposed on the surface. Across mammals, the relation between total cortical area A_t, thickness T, and exposed area A_e is

$$A_t \, T^{1/2} \propto A_e^{5/4}, \tag{12.1}$$

which indicates that the degree of folding increases with area in a systematic way and cannot be explained by the number of neurons alone [23]. The white matter that connects cortical regions grows faster than the gray matter, with a volume that scales approximately as the 1.23 power of gray matter volume [35]. Tension in the developing tissue has been proposed as a mechanism by which folding reduces the length of connections [32], and comparative studies relate cortical folding to brain size across primates [36, 19].

Worked Example 12.2 (Costs of doubling cortical area)**.** Suppose cortical area is doubled at the same thickness. The gray matter volume doubles to about 1000 cm^3. By the relation of [35], white matter volume rises by $2^{1.23} \approx 2.35$. If cortical metabolic power is proportional to gray matter volume and the cortex accounts for about half of the 20 W resting power of the brain, doubling the area adds about 10 W, which is about ten percent of whole-body resting power. Under Equation (12.1), the exposed area rises by $2^{4/5} \approx 1.74$, so the cortex is folded more deeply and the proportion of area lying within sulci increases. The head must accommodate the larger volume, which is a constraint on birth if the growth takes place before it and on postnatal skull expansion if it takes place after it.

12.3 The Fractal Brain Keel as a Proposal

The Fractal Brain Keel is a proposal that a developmental scaffold with a self-similar geometry could increase the area available for cortex during growth, in the way that branching geometries in plants and in the lung increase surface for a given volume. The proposal is to be evaluated through four separable questions, each with its own measurable quantity. The first is surface area, which is the quantity that fractal geometry increases. The second is wiring distance, which determines conduction delay and the volume of white matter, and which grows with area according to the scaling above. The third is transport, namely the delivery of oxygen and glucose and the removal of heat and waste, which in the cortex depends on the vascular supply and on cerebrospinal fluid flow. The fourth is coordination, the organization of activity across regions, which depends on the architecture of connections and not only on their number.

Worked Example 12.3 (Wiring delay under area doubling)**.** If the cortex is scaled in shape so that its area doubles, its linear dimensions rise by $\sqrt{2}$. A path of 10 cm across the hemisphere, conducted at 10 m s^{-1} in a moderately myelinated fiber, takes 10 ms, and the same path takes 14 ms after scaling if the conduction velocity is unchanged. Delays of this size are comparable to a substantial part of the period of a 40 Hz rhythm, which is 25 ms, and a thirty percent increase in cross-hemisphere delay therefore matters for any function that depends on timing. Raising the velocity to hold the delay fixed requires thicker or more

heavily myelinated fibers, which is a further cost in white matter volume and is consistent with the faster-than-linear growth of white matter reported in [35].

Each of these quantities can improve or worsen independently. A scaffold that increases surface area while lengthening wiring and demanding more transport could leave function no better than before, or worse.

Evidence Status

Established physiology. Cortical area, thickness, folding relations, and white-matter scaling are measured quantities across species [23, 35, 36, 19]. Sensory substitution and cortical remapping are established findings [3, 8]. Demonstrated engineering. No fractal developmental scaffold has been built. Proposed design. The Fractal Brain Keel is a proposal, and the discussion above specifies how it would be evaluated.

Exercise 12.1. An experiment is proposed in which cortical area is increased in an animal model by a developmental scaffold. State three outcome measures that would distinguish useful additional cortical organization from an increase in surface area alone, and one control condition that would separate the effects of the scaffold from the effects of larger brain volume.

Exercise 12.2. For the sensor of Worked Example 12.1, find the number of sites resolvable on the forearm if the threshold there is 3 cm and the area used is 0.1 m^2, and state which body site would give the highest capacity per unit of garment area.

Exercise 12.3. Using the exponent in Equation (12.1), find the exposed area when total cortical area is increased by a factor of four at constant thickness, and the fraction of total area that is hidden in sulci before and after the change, taking the initial fraction hidden to be two thirds.

Exercise 12.4. A sensor on the integument at the waist replaces a tactile sensor on the hand. Describe the training schedule and the outcomes that would show that the sensation has been remapped and not merely replaced by a cognitive translation.

Part IV

Three Architectures for a Redesigned Body

Chapter 13

The Hepastitium

The hepastitium is a proposed distributed auxiliary organ: a population of small implanted nodes joined by a relay network, intended to sample physiology, give local support to tissue, and recover fluid, in such a way that the loss of any single node is unimportant and the loss of many degrades the organ gradually [9, 10]. The proposal is treated here as a full case study, in the sense that every claim is stated as a specification that can be checked against the tools of Parts I and II. The aim is to find out what the design requires, and not to argue that it can be built.

13.1 Functions and Where They Live

Three functions are assigned to the organ, and they differ in what they need from placement and coordination. Physiological sampling measures analytes and physical quantities in the interstitial fluid close to tissue. Its requirement is proximity to the quantity sampled and a link by which the reading reaches wherever it is used. Local support acts on the tissue that a node serves, for example by modulating local perfusion, delivering a compound, or performing a small repair. Its requirement is proximity and a store of whatever it delivers. Fluid recovery moves water or a fluid fraction from a region where it is in excess to a place where it can be returned or discharged. Its requirement is a transport path, and the path ends either at the circulation or at an exit such as the Umbilicus of Chapter 14.

It follows that sampling and local support can be local to a node, and fluid recovery depends on the mesh. A further function, regulation, depends on the mesh in a different way, because it compares readings across nodes and adjusts set points. The division gives the specification of the organ in terms of the allocation matrix of Chapter 2: the column for the hepastitium has entries in the sensing, signaling, storing, sorting, and repairing rows, and those entries are distributed over the tissue volume.

Definition 13.1 (Organ at the level of the body)**.** A population of devices constitutes an organ at the level of the whole body if there is a function of the body, stated in the specification $\mathcal{S}$, that is carried out by the population as a whole, with a capacity that depends on the number and arrangement of working members and with a control that depends on the coordination among them.

The hepastitium satisfies the definition only if the mesh does real work, which means that the function of the whole depends on messages between nodes. A population of independent sensors reporting to an external device is a sensor array and not an organ.

13.2 Partial Support after Failure

If the organ is to retain partial support after failures, the geometry must permit it. The coverage relation of Equation (8.5) gives the fraction of tissue left unsupported when working nodes are scattered with density λ and each serves a sphere of radius r_c. A tissue region fails when no working node lies within r_c of it. The relation of Chapter 8 to the threshold model is as follows: the number of working nodes needed is determined by coverage, and the reliability of the whole is the probability that coverage and connectivity both hold.

Worked Example 13.1 (Sizing a regional organ)**.** Take $r_c = 2$ mm, so each node serves a volume of $\frac{4}{3}\pi r_c^3 \approx 33.5$ mm^3. To leave no more than one percent of the tissue unsupported under Equation (8.5), the dimensionless exponent must satisfy $\lambda V_c \geq \ln 100 \approx 4.6$, so $\lambda \approx 0.137$ mm^{-3}, or about 137 cm^{-3}. A liter of tissue then needs about 1.4×10^5 working nodes.

The power budget limits this. If each node draws 10 μW on average, a liter draws 1.4 W and the power density is 1.4 mW cm^{-3}, which lies below the heating bound of 15 mW cm^{-3} obtained from Equation (6.4) for a centimeter-scale implant. Covering the whole body, about 70 L, at this density would take about 10^7 nodes and 98 W, which is roughly the entire resting metabolic power of the person. The design is therefore either regional, covering only the tissues for which support is valuable, or run at a much lower average power per node by duty cycling, at the cost of the latency of sampling and support.

The example illustrates how a specification can be tested before any device exists. The question of what to cover is a question about the allocation of the organ, and the numbers show that complete coverage of the body at fine resolution is excluded by the energy budget and not by the geometry.

13.3 Mesh, Coordination, and Common Causes

Connectivity requires a mean number of neighbors above the percolation threshold quoted in Chapter 8. The communication range r_{com} must therefore be larger than r_c if connectivity is to hold at the density that coverage requires: at $\lambda = 0.137$ mm^{-3} and $r_{\mathrm{com}} = 2$ mm, the mean neighbor count is $\lambda \cdot \frac{4}{3}\pi r_{\mathrm{com}}^3 \approx 4.6$, which exceeds the threshold of about 2.7 with a margin of less than a factor of two. A localized loss of nodes, as from a tissue injury or a fibrotic capsule, can therefore break the mesh into islands that still cover their own tissue but no longer coordinate.

The common-cause term of Equation (8.4) is the other limit. The relay network needs power, whether harvested locally or delivered, and a protocol with shared state. If either fails for all nodes together, the organ fails regardless of its redundancy, and the design question becomes the reliability of that layer. For a distributed organ the common layer is the place where engineering effort should be concentrated.

Evidence Status

Established physiology. The tissue environment, interstitial compartments, and thermal limits for implants are standard physiology and physics [5, 18]. Demonstrated engineering. Individual implanted sensors, drug-delivery implants, and wireless power

links exist, as do ensembles of small devices in other settings, but there is no demonstrated body-scale network of implanted nodes. Proposed design. The hepastitium is a proposal developed in the cited manuscripts [9, 10], and this chapter provides analyses of its coverage, power, and connectivity requirements.

Exercise 13.1. For Worked Example 13.1, compute the power per node at which covering a 70 L body would draw 5 W. State how often a node could sample if each reading costs 1 μJ.

Exercise 13.2. Estimate the fraction of nodes that must be lost before the mean neighbor count falls to the percolation threshold, assuming uniform random loss. Compare this with the fraction that must be lost for the unsupported fraction in Equation (8.5) to exceed ten percent, and state which condition fails first.

Exercise 13.3. Specify what would have to be measured in an animal model to establish that the sampling function of a node population improves control of a physiological variable, giving the variable, the comparison condition, and the outcome.

Chapter 14

The Umbilicus

The Umbilicus is a proposed elective adult exchange interface, located at the navel, through which the body could be coupled to external metabolic infrastructure for nutrient delivery, waste removal, fluid regulation, and intermittent metabolic support [9]. As a design it is the opposite of the hepastitium in one respect: where the hepastitium moves function into the body, the Umbilicus lets function be performed by equipment outside it, and the body keeps only the interface. The central question is how much of the ordinary capacity of the body must be retained so that the person can disconnect, travel, or change infrastructure.

14.1 Three Modes of Coupling

Continuous coupling keeps the person connected for most of the day, and it places a premium on the reliability of the interface and of the infrastructure. Intermittent coupling connects the person for scheduled intervals, as in dialysis, and relies on the body to bridge the gaps. Emergency coupling is rarely used and demands that the interface remain usable after long periods of inactivity. The modes differ in their grace periods, in the burden of maintaining the interface, and in the amount of native function that must be preserved.

Table 14.1: Coupling modes compared on the quantities of Chapter 4.

Property	Continuous	Intermittent	Emergency
Native capacity required	Low in normal use, full on disconnection	Enough to bridge gaps between sessions	Full
Grace period needed	Hours for equipment faults	Length of the longest gap	Not applicable until used
Interface burden	Highest; a permanent portal	Moderate; may be closed between sessions	Lowest, but must work when first used
Main hazard	Infection and mechanical failure at the portal	Cycle of access and healing	Disuse of the interface

14.2 What Must Remain in the Body

Suppose that an external system takes over a function with load D (in the units of that function) for a fraction u of each cycle of length T_c, and that the body must supply the function alone in the remaining time. If the function accumulates a quantity in the body when unsupported, the stock reaches its peak at the end of the gap,

$$S_{\max} = (\dot{s}_{\text{gen}} - C_{\text{nat}})\,(1 - u)\,T_c, \tag{14.1}$$

where $\dot{s}_{\mathrm{gen}}$ is the rate of generation or load and C_{nat} the native capacity. If the stock must remain within the viable limit S_{lim}, then the native capacity must be at least

$$C_{\mathrm{nat}} \geq \dot{s}_{\mathrm{gen}} - \frac{S_{\mathrm{lim}}}{(1-u)\,T_c}. \tag{14.2}$$

The capacity that is retained therefore trades against the length of the gap. A person who wants to be able to leave for a long time must retain more native function, and a person who accepts frequent coupling can retain less.

Worked Example 14.1 (Urea and the length of a disconnection)**.** An adult produces about 0.5 mol of urea per day, distributed through about 42 L of body water. With no clearance at all, the concentration rises at $0.5/42 \approx 12$ mmol L^{-1} per day. Taking a baseline of 5 mmol L^{-1} and a limit, for the purpose of the example, of 40 mmol L^{-1}, the rise permitted is 35 mmol L^{-1}, which at 12 mmol L^{-1} per day gives a grace period for this function of about three days. If the person retains native clearance of one third of the generation rate, the net rise is two thirds as large, and the grace period lengthens to about four and a half days. If the person retains half, it is six days. The retained native function buys time for the other functions to be arranged, and the relation is hyperbolic: the grace period diverges as native clearance approaches generation.

The analysis shows that the specification of the Umbilicus is not complete without stating the retained function. A design that gives up the kidney entirely has a grace period in days for urea and in a smaller number of days for potassium and acid, while a design that retains partial kidney function has a longer one. The same reasoning applies to nutrition, where the body's stores give a grace period in weeks, and to water, where it is hours to days depending on the environment.

14.3 The Interface Itself

A permanent opening through the skin is a standing breach in the barrier that the epithelium maintains. Percutaneous drivelines of ventricular assist devices are a recognized site of infection and are among the main complications of those devices [22]. The design has to say how the interface is sealed, how the tissue at its margin is kept healthy, how the interface is cleaned, and whether it can be closed when the person is not coupled. The navel is a plausible location because the abdominal wall there is thin and has a history of closure after development, and because it is on the trunk, where it is covered by clothing and is less exposed to impact than the limbs. These advantages do not remove the problem of infection, and the evaluation of the interface must treat it as the principal risk.

Evidence Status

Established physiology. Urea production, body water, and the course of unsupported renal failure are standard clinical physiology [5, 18]. Demonstrated engineering. Dialysis access, parenteral nutrition, and percutaneous device drivelines are in clinical use with documented complications [22, 17]. Proposed design. The Umbilicus as an elective interface for healthy adults is a proposal [9].

Exercise 14.1. Using Equation (14.2), find the native capacity, as a fraction of the generation rate, that allows a person coupled for 8 hours of every 24 to remain below the limit of Worked Example 14.1. Repeat for a person coupled for 2 hours every 48.

Exercise 14.2. Write the specification of an emergency coupling in the notation of Chapter 4, giving the disturbance class, the viable set, and the requirement on the interface after one year of disuse.

Chapter 15

The Electric Jacket Becomes Skin

Garments already do a part of the work of the body. Heated jackets manage warmth, instrumented garments sense heart rate and movement, and padded garments reduce impact. The proposal of this chapter is a continuous path from such a garment to a permanent integument, with a specific next step: once water leaves the body, the garment retains it inside a managed system from which it can be collected, treated, or discharged later.

15.1 The Garment as a Boundary Layer

A garment is a stack of layers in the sense of Chapter 3. The innermost layer contacts the skin and manages moisture and comfort. A fluid-handling layer moves liquid by capillary action or by pumping toward a reservoir. A thermal layer provides insulation or active heating or cooling. A structural layer carries load and absorbs impact. A layer of power, sensing, and communication ties the rest together, and an outer shell protects the stack from the environment. Each layer has a characteristic life, and a design that is meant to be permanent has to assign each an interval of renewal.

Worked Example 15.1 (The limit of battery-powered warmth)**.** A heated jacket that supplies 30 W over an eight-hour day uses 240 Wh. Lithium-ion cells with a specific energy of about 200 $\mathrm{Wh\,kg^{-1}}$ would take 1.2 kg of cell mass to carry this energy, which is practical. To supply 100 W, the order of the resting metabolic output of the person, for eight hours would take 800 Wh and 4 kg of cells, which is a burden but not an impossible one. The energy has to be delivered at the skin, where the conversion of electric power to heat is nearly complete, so the cost is set by the energy and the battery mass and not by efficiency.

The same energy, if used to power a vapor compression cycle for water recovery at 100 $\mathrm{kJ\,kg^{-1}}$, would handle 1.5 kg of water for 150 kJ, which is 42 Wh, much less than the heating demand. The comparison shows that a worn system's energy budget is dominated by thermal management and that fluid handling is a small increment on it. A permanent integument should therefore minimize the active power spent on temperature, by relying on insulation and on control of evaporation, and reserve the electrical budget for sensing and fluid handling.

15.2 From Garment to Integument

A permanent integument is a layered system that is worn continuously and whose components are renewed routinely. The renewal strategy can be stated by assigning each layer an interval: the contact layer, which accumulates skin cells, salts, and microorganisms, is renewed in days; the fluid-handling layer, which fouls with solutes, in weeks; the thermal and structural layers in years; and the outer shell as wear requires. Renewal can be done at a service point,

where used layers are removed, cleaned, and remanufactured, and where the stored fluid is processed. The person's skin lies beneath the contact layer and has to remain healthy. Continuous occlusion of skin leads to maceration, which is the softening of the outer layer by retained moisture, so the contact layer must pass water vapor and drain liquid at rates that match the person's sweating, which in exertion exceeds 1 $\mathrm{L\,h^{-1}}$.

The step from garment to integument is also a step in how the person relates to the layer. A removable garment can be taken off, and a permanent integument cannot be without a procedure. The specification of Chapter 4 then has to include the requirement that the person survive a failure of the integument, which could happen far from a service point. The grace period for a failure of the fluid-handling layer depends on how much water it holds and on whether it can be vented, and for a failure of the thermal layer it depends on the environment.

15.3 Fluid Handling as the Distinctive Function

The distinctive contribution of the integument is that it treats the body's fluid losses as a managed stream. Sweat, insensible loss, and, in some designs, other fluids are collected in the layer, and the water is separated and stored, treated, or released as vapor according to the control law of Chapter 7. The Worked Example 7.1 shows that the control law must give priority to heat rejection when the environment is hot, because the outer surface of the garment is a single budget for both heat and retained water. In cold or temperate conditions, the integument can retain nearly all fluid and the water balance of the person improves, and in heat it must open an evaporative path and accept losses.

15.4 Regulation by Geometry: The Thermovariable Weave

The control problem of Chapter 7 asks for a layer that changes its permeability with the thermal state of the wearer. One way to build such a layer uses no electronics. Strips of bimetal, in which two metals of different thermal expansion are bonded together, bend when their temperature changes, and a weave that contains such strips opens or closes its pores as the fabric warms or cools. The textile then acts as its own actuator and its own sensor, and its geometry makes the state of the garment visible to the wearer [11]. The idea has been extended in design discussions to resistive heating and to fluid-circulating microtubing, and to a relay mesh that coordinates the response of separate patches. These extensions are design concepts, and no prototype or performance measurement of the combined system is reported.

For a bonded strip of two layers of equal thickness and equal stiffness, the curvature produced by a temperature change ΔT is approximately

$$\kappa \approx \frac{3\,\Delta\alpha\,\Delta T}{2h}, \tag{15.1}$$

where $\Delta\alpha$ is the difference between the expansion coefficients of the two layers and h is the total thickness [30]. A cantilevered length ℓ then deflects at its tip by $\delta \approx \kappa\ell^2/2$.

Worked Example 15.2 (Opening a pore)**.** Take layers of aluminum and steel, with $\Delta\alpha \approx 11 \times 10^{-6}\ \mathrm{K^{-1}}$, a total thickness of $h = 100\ \mu\mathrm{m}$, and a free length of $\ell = 20$ mm. For

a temperature change of 15 K, Equation (15.1) gives $\kappa \approx 2.5\ \mathrm{m}^{-1}$ and a tip deflection of $\delta \approx 0.5$ mm. This is of the order of the gap between yarns in an open textile, so a weave of such strips can plausibly change its openness by a useful factor over a range of fabric temperature of the size that occurs between a cool and a warm skin surface. A strip of a quarter of the length deflects sixteen times less, so the design favors long, thin strips, and the weave geometry has to accommodate the free length.

The device differs from the controller of Chapter 7 in three ways. Its set point is fixed by the materials and the geometry and cannot be adjusted for the person's state, whereas the native controller reads core temperature and the rate of heat storage. Its signal is the temperature of the fabric, which is a mixture of skin and ambient conditions, so it inherits the problem of disagreeing sensors. Its delay is the thermal time constant of the strip, which for a thin strip is seconds, far shorter than the delay of the human controller. The weave is therefore well suited to the fast, local part of regulation, and the decision to open the evaporative path in response to heat storage still requires a slower loop that sees the body.

Exercise 15.1. For the strip of Worked Example 15.2, find the length needed for a 1 mm deflection at a temperature change of 10 K, and the number of cycles of bending the strip undergoes in a year of daily use. Discuss what limits the life of the weave.

Evidence Status

Established physiology. The structure of skin, sweating rates, and the physiology of maceration are standard [18, 29]. Demonstrated engineering. Heated and instrumented garments, wicking fabrics, and liquid-cooled suits are in use [7]. The bimetal relation of Equation (15.1) is a standard result [30]. Proposed design. The permanent integument with routine renewal and managed fluid retention, and the thermovariable weave with its extensions [11], are design concepts.

Exercise 15.2. For a garment of areal density $1\ \mathrm{kg\,m}^{-2}$ over $1.8\ \mathrm{m}^2$, find the mass of the garment and the mass of a fluid-handling layer that can hold 1 L of water at ten times its dry mass. Discuss which of these dominates the burden to the wearer.

Exercise 15.3. Assign renewal intervals to the six layers described in this chapter and compute the number of layer changes in a year. State the consumable and the waste stream produced by each.

Chapter 16

The House, Car, Fridge, Stove, and Sink

Ordinary life depends on fixed equipment that stands in for functions the body does not have. A house supplies shelter and a habitable climate. A car extends mobility and carries loads. A refrigerator and a stove alter the timing and availability of food. A sink supplies and drains water. This chapter translates each appliance into the service it provides and asks which services could share hardware on the body, which benefit from a wearable buffer, and which remain more practical at the scale of a building or a shared utility.

16.1 Services and Their Scales

Table 16.1: Appliance services and the scale of their loads for one person. Values are order-of-magnitude figures.

Appliance	Service	Scale of load	Natural placement
House	Climate and shelter	Kilowatts of heat exchange for a dwelling, against about 100 W of body heat	Building, with a worn microclimate layer
Car	Mobility and load carrying	Tens of kilowatts at speed, against a walking cost of about 2 to 3 $\mathrm{J\,kg^{-1}\,m^{-1}}$	Vehicle
Fridge	Delay of spoilage	Tens of watts of cooling for a household unit	Building or shared utility
Stove	Heating of food	About 1 kWh per person per day for domestic cooking	Building; small worn heater possible
Sink	Water supply and drainage	Tens to hundreds of liters per day, against about 2.5 L drunk	Building and utility

The table shows that the services differ by orders of magnitude in the scale of load, and the scale suggests where each can sit. The fixed climate of a house is far larger than the body's heat output, and no worn system can replace a house in cold weather, but a worn layer can extend the range over which a person is comfortable and thereby reduce the load that the house must meet. The water supply of a sink is, by volume, mostly for hygiene and cleaning, and the quantity that the body needs to drink is a small part of it. A worn recovery system addresses the small part and leaves the large part to the building.

16.2 Sharing Hardware

Several services have hardware in common. Heat is the common currency: the stove adds heat to food, the fridge removes it, the house adds or removes it from air, and the body

needs it or rejects it. A worn thermal store could supply heat for warming the body in the cold and for warming food, and could receive heat rejected in recovering water. Power is another shared resource, and a battery that powers a garment can also power a small cooker or a lamp. Water handling is shared between the integument and any appliance that uses water, because both need filtration and storage.

Worked Example 16.1 (Heating food from a worn store)**.** To warm 1.5 kg of food per day from 5 °C to 60 °C with a specific heat of 3.5 $\mathrm{kJ\,kg^{-1}\,K^{-1}}$ requires $1.5 \times 3.5 \times 55 \approx 290$ kJ, or 80 Wh. A 100 Wh battery covers this with a margin, at a mass of 0.5 kg. The warming of food is thus well within the reach of a worn energy store, as long as the store is not also used for body heating. Cooking in the sense of raising food to boiling and keeping it there takes more, and the domestic figure of about 1 kWh per day in Table 16.1 is a factor of twelve larger, mostly because the stove heats the surrounding vessel and air and because much of the food is cooked in larger quantities than one person eats. A worn store is therefore enough for reheating and not for a full domestic kitchen.

16.3 The Design Problem

The design problem of the chapter is to provide a day of ordinary life with less dependence on fixed appliances. Write the day as a sequence of services with their loads and durations, assign each to one of three places, the body and its worn layers, a vehicle or portable unit, or a fixed utility, and compute the carried mass and the fixed dependence that results. A design is better if it reduces fixed dependence without raising carried mass beyond a limit, and if it preserves the grace period for each service. The three-way allocation is the precursor of the accounts of Chapter 17.

Evidence Status

Established physiology. The values for body heat production, food energy, and the specific heat of food are standard [18], and the cost of walking is a standard biomechanical value. Demonstrated engineering. Portable stoves, heated garments, and personal water treatment devices are in common use. Proposed design. The sharing of hardware among these services on the body is a design concept.

Exercise 16.1. Compose a table of one day in which each service is assigned to the body, a portable unit, or a fixed utility. Compute the mass carried and the fraction of each service's load that is supplied from the fixed utility.

Exercise 16.2. Identify a service in Table 16.1 for which sharing hardware on the body would increase the risk of simultaneous failure, using the common-cause relation of Equation (8.4), and propose a way to reduce that risk.

Part V

Smaller Bodies, Larger Operating Ranges

Chapter 17

Functional Compaction

The conjecture that a person could carry less material while living and acting in a wider range of conditions is stated in this chapter as an engineering problem. The statement has to say which functions are to be kept, in which environments, and what counts as carried material. A conjecture that is stated this way can be tested, and it can fail in identifiable places, which is the subject of Chapter 18.

17.1 Four Accounts

Carried material is divided into four accounts that are kept separate in every calculation. Biological tissue, with mass m_b, is the body that the person is born with, modified by any engineered additions to it. Worn hardware, with mass m_h, is equipment on the body that is not part of it, including garments, interfaces, and packs. Consumables, with mass $m_c(t)$, are materials that are used up or accumulate during the period of use, such as water, food, filter media, and brine. Services supplied at a distance are not carried, and their cost is recorded as a dependence and not as a mass.

The need for separate accounts is that a design can lower one at the expense of another and so appear to reduce the total, as when a person is lightened by moving a function into a service that must then be available. Mass carried is a measure of burden, and dependence is a measure of exposure.

17.2 The Problem

Let A denote an architecture, meaning an assignment of each function in a specification to the body, worn hardware, or a distant service, together with the devices that realize it. Let $\mathcal{S} = (\mathcal{K}, X_0, T, \text{task})$ be a functional specification as in Definition 4.1. The compaction problem is

$$\min_{A} \; m(A) = m_b(A) + m_h(A) + m_c(A;T) \quad \text{subject to} \quad A \,\text{satisfies}\, \mathcal{S}, \;\; \rho_{\text{res}} \geq \rho_{\min}, \;\; \tau_g \geq \tau_{\min}, \tag{17.1}$$

with a second objective, the exposure vector $E(A)$, whose components are the fractions of each service supplied at a distance. The solution is a Pareto set in the plane of carried mass and exposure, since lowering the first generally raises the second. The operating range $R(A)$ of an architecture is the set of environments in which $\mathcal{S}$ is satisfied. Two distinct conjectures are in play, and either can hold without the other. The first, the carried-mass conjecture, is that there exists an architecture A with $m(A) < m(A_{\text{ref}})$ and $R(A) \supseteq R(A_{\text{ref}})$, where A_{ref} is a reference with the ordinary anatomy and standard equipment. The second, the anatomy conjecture, is that a body with a smaller biological tissue mass m_b preserves the specification, which is a statement about the body and not about what it carries. Each

is stated for a declared reference anatomy, reference equipment, set of environments, and time horizon. Consumables $m_c(A;T)$ mean the peak inventory over the journey of usable materials, waste that is carried, and replaceable media, and implanted hardware is placed in the worn-hardware account m_h, reported separately as m_h^{imp} so that it can be distinguished from external hardware. A proof of either conjecture requires one feasible architecture whose complete ledger and joint operating range meet the comparison.

Three families of architecture are compared. Shared infrastructure inside the body reduces m_h by placing hardware in tissue, and it is limited by tissue heating and by the difficulty of maintenance. Multifunctional materials at the surface reduce m_h by letting one layer perform several services, as in the integument of Chapter 15. Access to external equipment reduces all carried accounts and increases E, as in the Umbilicus of Chapter 14.

Worked Example 17.1 (Water for a multi-day journey)**.** A person walks in a temperate climate for T days and needs a throughput of 4 kg of water per day. In the reference architecture all water is carried at the start, so the peak carried mass is $m_{\mathrm{ref}} = 4T$ kilograms.

In the recovery architecture, let f be the fraction of the throughput that reaches a collection stage and η the recovery fraction of the collected stream. With a steady inventory and no other water inputs, the daily makeup is $4(1 - f\eta)$ kilograms. The water that is not recovered leaves the usable inventory and enters the waste inventory at the same moment, so it must not be counted twice. Take as the waste policy that the brine, at 0.2 kg per day from Worked Example 5.2, is stored, and that the remaining unrecovered water is discharged along the route. For $f = 1$ and $\eta = 0.9$ the net loss of usable water is 0.4 kg per day, of which 0.2 kg moves to the brine store and 0.2 kg is discharged.

The usable inventory at the start must cover one day of throughput plus the net losses of the whole journey, $4 + 0.4T$ kilograms. The unit has a mass of 1.5 kg, the replaceable media consumed amount to 0.13 kg per day, and the energy of 42 Wh per day takes cells of 0.22 kg per day at 200 $\mathrm{Wh\,kg^{-1}}$. The brine store is empty at the start and grows by 0.2 kg per day while the usable inventory falls by 0.4 kg per day, so the carried mass falls over the journey and its peak is at the start,

$$m_{\mathrm{peak}} = 1.5 + 4 + \big[4(1 - f\eta) + 0.13 + 0.22\big]T = 5.5 + 0.75\,T \quad (f = 1,\ \eta = 0.9).$$

The break-even against $4T$ is $T^* = 5.5/(4f\eta - 0.35) \approx 1.7$ days. For $T = 3$ the peak carried mass is 7.75 kg against 12 kg, and for $T = 10$ it is 13 kg against 40 kg. If only 80 percent of the throughput is collected, as when sweat in hot conditions is not captured, then $f\eta = 0.72$ and $T^* \approx 2.2$ days. The result is conditional on the collection fraction, the heat load, the waste policy, and the resupply interval, and it applies to the environment e_1 of the running problem. In e_2 the garment must open its evaporative path, f falls, and the break-even moves out.

17.3 Reading the Accounts

The example illustrates the use of the four accounts. The saving is entirely in the consumables account, and the other accounts rose: worn hardware by 1.5 kg for the unit, and dependence by the need for replacement media. The operating range is narrower on the hot side unless a further provision is made, so the carried-mass conjecture is not yet established by this

architecture, and a full design would add the provision and recompute. The anatomy conjecture is untouched by the example, since no tissue was changed. The structure of the argument, which is to state accounts, compute the break-even, and test against the range, is the method of the chapter.

17.4 A Pareto Example

The break-even calculation is a one-dimensional comparison. The problem of Equation (17.1) has two objectives, and its solution is a set. Because an architecture that is lighter but has a narrower operating range has not met the comparison, dominance is defined with range containment: A dominates A' only if $m(A) \le m(A')$, $E(A) \le E(A')$, with one strict, and $R(A) \supseteq R(A')$. Suppose six architectures for the three-day journey of Worked Example 17.1 have been analyzed, with carried mass in kilograms and exposure as the fraction of services supplied at a distance.

Worked Example 17.2 (Finding the Pareto set)**.** The six pairs (m, E) are as follows, and the values are illustrative. For simplicity every architecture is assumed to preserve the reference operating range, so that range containment holds among all of them, and an exercise removes this assumption. The reference architecture, which carries all its water, has $(12.0, 0)$. The worn recovery unit of Worked Example 17.1 has $(7.75, 0.10)$, the small exposure being the resupply of filter media. A variant in which the recovery function is built into a shared integument has $(9.0, 0.10)$. A scheme that relies on caches of water along the route has $(6.5, 0.35)$, and a hybrid of recovery and one cache has $(6.0, 0.40)$. A scheme that is coupled to a vehicle or a utility for most services has $(4.0, 0.80)$. Finally a poorly designed variant has $(7.0, 0.45)$.

An architecture is dominated if another has no larger mass and no larger exposure, with one of the two strictly smaller. The integument variant is dominated by the worn recovery unit, which has the same exposure and a smaller mass, and the poorly designed variant is dominated by the hybrid, which is lighter and less exposed. The Pareto set consists of the reference, the worn recovery unit, the cache scheme, the hybrid, and the coupled scheme. Among these none is best without a statement of how mass is weighed against exposure, and moving along the set trades about 4.25 kg of mass for 0.10 of exposure between the first two, and about 1.25 kg for 0.25 between the second and the third, so that the marginal price of exposure differs along the set. The step to the coupled scheme buys 2.0 kg for 0.40 of additional exposure and is the most expensive in dependence.

The example shows the use of the Pareto framing: it removes the dominated designs without any weighting of the objectives, and it exposes the marginal rates at which one objective is traded for the other, which is the information a person needs to decide how much exposure they are willing to accept.

Evidence Status

Established physiology. The accounts use the physiological quantities of Part II. Demonstrated engineering. Portable water treatment and recovery devices exist and their masses and consumables are documented [7]. Proposed design. The smaller-body conjecture and its formalization as Equation (17.1) are proposals developed in this book.

Exercise 17.1. Repeat Worked Example 17.1 for a recovery fraction of 0.8 and for 0.95. State how the break-even time changes, and identify the account that dominates the slope in each case.

Exercise 17.2. Add to Worked Example 17.1 a hot-climate day in which the unit is bypassed for six hours and the person loses 2 L of sweat to evaporation. Recompute the carried mass and the break-even time, and determine whether the architecture still satisfies the specification at 35 °C.

Exercise 17.3. Suppose the worn recovery unit loses the hot half of the reference operating range. State which architectures it then no longer dominates, and show that the Pareto set under range containment can contain a heavier architecture that the mass-and-exposure relation alone would discard.

Exercise 17.4. Add to Worked Example 17.2 a seventh architecture with $(5.0, 0.30)$ and determine which members of the Pareto set it dominates.

Exercise 17.5. Define a second Pareto plane for the architectures of this chapter, with grace period on one axis and carried mass on the other. Argue which of the three families lies furthest from the origin of the plane.

Chapter 18

Where Miniaturization Stops Paying

A smaller system has less room for reserves, must reject heat across less area, and is harder to maintain. Shrinking a component helps until one of these effects dominates, and the point at which the effect dominates is where compaction stops paying. This chapter examines the effects, compares shrinking a component with shrinking the whole system around it, and gives the central exercise of tracing where an apparent volume saving went.

18.1 Heat at Fixed Function

A component that must deliver a fixed power P does not dissipate less power when it is smaller. Its power density rises as its volume falls. For a spherical implant in tissue, substituting $p_d = P/(\frac{4}{3}\pi R^3)$ into Equation (6.4) gives

$$\Delta T = \frac{P}{4\pi k R}. \tag{18.1}$$

At fixed power the temperature rise grows as the radius falls. For $k = 0.5\ \mathrm{W\,m^{-1}\,K^{-1}}$, a unit dissipating 60 mW with $R = 1$ cm rises by 0.95 K, and one dissipating the same power with $R = 5$ mm rises by 1.9 K. A unit dissipating 1 W with $R = 1$ cm rises by 16 K, which is not tolerable without perfusion cooling. Miniaturization that is accompanied by constant function therefore has a floor set by tissue heat removal, and the floor rises with the power that the function needs.

18.2 Reserves and Endurance

The scaling of Equation (8.2) shows that endurance declines with size for a fixed fractional reserve. A design that shrinks reserves in order to save carried mass lowers the grace period in proportion. If the reserve of a consumable is S and its use rate is q, the grace period for a loss of supply is $\tau_g = S/q$, and a smaller reserve gives a shorter one.

Worked Example 18.1 (Where the saving went)**.** Two bodies have identical nominal performance in the sense that each meets the specification of Chapter 4 in the environments tested, with supply available. Body A carries a water reserve of 7.8 L, three days of the use rate $q = 2.6\ \mathrm{L\,day^{-1}}$, and Body B carries 1.3 L, half a day. Body B carries 6.5 kg less, and its apparent volume saving is 6.5 L. When a supply line is cut, Body A has a grace period of three days and Body B of half a day. The saving did not disappear. It was exchanged for the grace period, which was reduced by a factor of six, and for a standing requirement that the supply line be present.

The accounting identifies the saving as a transfer from the consumables account to the exposure vector E of Chapter 17. A design that reports the lower carried mass without stating the new exposure has hidden the transfer.

18.3 Shrinking the Component and Shrinking the System

Two cases of miniaturization are different. Shrinking a component, such as a pump or a membrane module, leaves the rest of the system unchanged, and the gain is the mass of the component, subject to the heat limit of Equation (18.1). Shrinking the whole system, such as the body itself, changes all the scaling relations at once: surface area falls as the square of the dimension, reserves as the cube, and basal power as a power of the mass between them, as in Equation (8.2). The result for the whole system is that the heat to be rejected per unit area changes, the reserve falls faster than the surface available to exchange with the environment, and the margins that were adequate for the larger body are not adequate for the smaller one unless the architecture changes. A uniform shrinking is therefore a different design problem from a component upgrade, and the conjecture of Chapter 17 requires a different analysis for each.

18.4 Maintenance at Small Scale

Small devices are harder to maintain because access is limited and because each failure is a smaller event that is harder to detect. A population of small nodes can tolerate failures, as the threshold model of Chapter 8 shows, but only if their failures are detected and the failed nodes are removed or bypassed, and the cost of detection is part of the maintenance account. Miniaturization shifts cost from the components to the process of maintaining them, and the shift can be larger than the saving.

Evidence Status

Established physiology. Heat conduction in tissue, scaling of reserves with size, and the structure of metabolism are established [21, 34, 4]. Demonstrated engineering. The thermal limits of implants are respected in the design of existing devices. Proposed design. The comparisons of this chapter are analyses of proposed architectures, and the accounting of where savings go is a method proposed in this book.

Exercise 18.1. Two bodies have identical nominal performance and different endurance when a supply line is cut. Using Worked Example 18.1 as a model, construct a different pair for oxygen or for electrical energy, state the reserves, and show exactly where the apparent volume savings went in terms of the accounts of Chapter 17.

Exercise 18.2. Using Equation (18.1), find the smallest radius of a spherical implant that dissipates 200 mW without a temperature rise above 2 K. Discuss how local blood perfusion, which adds a cooling term, would shift the result.

Chapter 19

Development Into a Different Anatomy

Three kinds of change can be compared. A jacket worn by an adult can be taken off. A permanent interface adopted by an adult is added to an anatomy that has already formed. An integument around which a person develops from infancy is part of the conditions in which the body forms. The questions of this chapter are how movement and mechanical surroundings shape development, and what a person must learn anew when the location of sensation, support, or bodily control is changed.

19.1 Mechanics in Development

Embryonic and fetal movement is needed for normal formation of joints and bones in several animal models, and paralysis during development gives abnormal joint shapes, reduced bone size, and altered muscle [25, 24]. The conclusion that mechanical conditions are inputs to development is well supported for the skeleton. The more specific claims that the structure of the extracellular matrix, and collagen in particular, records the mechanical history of development and that movement in the womb shapes sensorimotor organization are the subject of two research programs of the author, Mechanical Inheritance and Motile Womb Theory [12, 13, 14]. The earliest dated mention of the second idea located so far is a commit of December 2022 [14], and the developed account is in the later manuscripts. In this book they provide a framing: if surrounding mechanical conditions help shape the formation of the body, then an integument that supplies different mechanical conditions from infancy would be a change in the conditions of development and not only an addition to an adult.

19.2 What the Framing Adds

The standard literature on developmental mechanics already supports the statement that mechanical conditions influence skeletal development [25, 24]. The framing of Mechanical Inheritance adds claims that the standard account does not make, and that can be tested. The first is that the mechanical history of development is retained in the extracellular matrix, with a persistence set by the turnover of its components, so that the retention of developmental loading into adulthood should vary across tissues in the manner of Equation (24.1). The second is that the organization of sensorimotor function reflects the movements made during development in the womb, which Motile Womb Theory proposes and which would be tested by altering fetal movement in an animal model and measuring later sensorimotor organization. The third, specific to this book, is that a developmental integument supplies a different mechanical history and so should leave a trace in slowly turning tissues, which gives a measurement: comparison of matrix properties, such as the crosslinking and orientation of collagen, in the slow-turnover tissues of people raised with and without the integument. Without these claims the framing would add nothing to the

standard citations, and with them it makes predictions that could fail.

19.3 Growth and the Renewal Schedule

A person who grows within an integument needs one that grows with them. The routine renewal described in Chapter 15 offers a way to do so, because each renewed layer can be fitted to the current size.

Worked Example 19.1 (Growth between renewals)**.** A child between the ages of five and ten grows in height by about 6 cm per year from a height near 110 cm, a linear growth of about 5.5 percent per year. Surface area scales as the square of the linear dimension, so the area grows by about 11 percent per year. A layer renewed annually is, on average, undersized by half of that at the time of its replacement, about 6 percent in area, and is undersized by 11 percent at its end. An elastic layer that tolerates a strain of 20 to 30 percent can absorb this without refitting, and a layer that does not tolerate this much strain must be renewed more often. The surface area of the body grows by a factor of about seven between birth and adulthood, so across childhood the renewal schedule has to include some changes in the pattern, and not only the size, of the layers.

19.4 Learning a New Location of Sensation and Control

When a sensor is moved to a new site, the sensory map has to be rebuilt. Experiments with tactile substitution show that users learn to attribute sensations to external objects after training [3], and studies of amputees show that the cortical representation of the body changes after loss [8]. A person who develops with the integument has the advantage that the mapping can form together with the body, whereas an adult has to retrain. The corresponding cost is that a person who develops with a particular layout has no experience of the alternative and cannot judge what has been given up.

19.5 Questions for a Developmental Specification

A specification for development has to include, besides the functional specification of Chapter 4, constraints on the course of growth. Joint motion has to remain within the range for normal skeletal development, since immobilization alters bone shape. The thermal and fluid conditions at the skin have to remain within the range for normal skin development, and sensory input in the first years has to be sufficient for the formation of maps. Each constraint is measurable, and none is satisfied by default.

Evidence Status

Established physiology. Mechanical conditions influence skeletal development in animal models, and sensory and cortical remapping occur in adults [25, 24, 3, 8]. Demonstrated engineering. Growth-accommodating garments and prostheses for children exist. Proposed design. Mechanical Inheritance, Motile Womb Theory, and the application to a developmental integument are proposals [12, 13].

Exercise 19.1. Choose one of the three claims in the section above and design an animal experiment to test it, stating the outcome that would count against it.

Exercise 19.2. Recompute Worked Example 19.1 for an adolescent growth spurt of 9 cm per year from a height of 150 cm, and find the renewal interval that keeps the layer within a strain of 20 percent.

Exercise 19.3. Design an observation that would test whether movement during development shapes sensorimotor maps, and state what a negative result would imply for the proposal that a developmental integument changes organization.

Chapter 20

Maintenance, Dependence, and the Ability to Leave

A succession of individually useful interventions can alter the body's own capacity for repair and make future support harder to discontinue. This is the formal argument of the manuscript on continuous microsurgery [9], and this chapter gives students explicit ways to represent the loss of options: a grace period after support fails, the availability of independent maintenance, and the possibility of returning to a simpler configuration.

20.1 A Model of Eroding Native Capacity

Let the body require a repair rate r, normalized to one, to remain in the viable set, and let its native repair capacity be c, with $c_0 > 1$ initially. Each intervention i substitutes for some native repair and reduces native capacity by a fraction δ_i,

$$c_{i+1} = (1 - \delta_i)\, c_i. \tag{20.1}$$

Each intervention is beneficial in the sense that, with support present, damage is lower than before. If support is withdrawn while $c < r$, unrepaired damage accumulates at the rate $r - c$ until it reaches the limit $D_{\lim}$, and the grace period, for damage D_0 already present when support is withdrawn, is

$$\tau_g(c) = \frac{D_{\lim} - D_0}{r - c} \quad (c < r), \qquad \tau_g = \infty \quad (c \geq r), \tag{20.2}$$

with $\tau_g = 0$ if $D_0 \geq D_{\lim}$. Where recovery during withdrawal is possible at a rate ρ, the rate $r - c$ is replaced by $r - c - \rho$, and the grace period is infinite if that is not positive. In the examples below $D_0 = 0$, so the earlier expression $D_{\lim}/(r - c)$ applies. The approval of an intervention is governed by a rule, here taken as an assumption, that it is accepted whenever damage is lower with support present; Worked Example 20.1 is a counterexample to that rule as a sufficient criterion. The grace period is infinite while the person remains self-sufficient, falls steeply as capacity drops below demand, and approaches $D_{\lim}/r$ as capacity falls to zero.

Worked Example 20.1 (Ten small steps)**.** Take $c_0 = 1.5$ and $\delta_i = 0.10$ for each intervention. Native capacity falls below demand when $(0.9)^n < 1/1.5$, that is when $n > \ln 1.5/(-\ln 0.9) \approx 3.85$, so after four interventions the person is no longer self-sufficient. Let $D_{\lim}$ be the damage equal to six months of the full repair demand. After five interventions $c_5 = 1.5 \times 0.9^5 \approx 0.886$ and $\tau_g = 6/0.114 \approx 53$ months, about four and a half years. After ten, $c_{10} = 1.5 \times 0.9^{10} \approx 0.523$ and $\tau_g = 6/0.477 \approx 12.6$ months. Each step was beneficial and cost a tenth of the native capacity, and the sum of ten steps took the grace period from infinite to about a year. If the steps are smaller, say $\delta = 0.05$, it takes eight to cross the threshold, and the shape is the same.

No single step in the sequence was dangerous in itself, and the loss of options is a property of the sequence. A design review that examines each intervention on its own will approve every one.

20.2 Three Representations of the Loss of Options

The first representation is the grace period of Equation (20.2), a number of days or months that gives a direct measure of how long the person can be without support. The second is the availability of independent maintenance, a number $a \in [0, 1]$ equal to the fraction of maintenance tasks, weighted by their importance to the specification, that the person can perform without specialized service. An integument that can be renewed by the wearer has high a, and an implant that requires a surgeon has low a. The third is reversibility, a property of the state: a configuration is reversible if there is a sequence of configurations, each of which satisfies the specification, that leads to a designated simpler configuration. Reversibility depends on the capacity to rebuild what was lost, so it is limited by the rate ρ_{rec} at which native capacity can be restored. If restoring takes longer than the grace period, the return path is closed even though the final configuration would be viable.

Definition 20.1 (Ability to leave)**.** Let a configuration x be given with native capacity $c(x)$ and support u. The person has the ability to leave at x if, for the best available fallback, there exists a schedule of reduction of support such that the state remains in $\mathcal{K}$ throughout and ends in a configuration that requires no service beyond what the person can supply. The ability to leave is graded by the minimum grace period over the schedule and by the independent maintenance availability a at its end.

20.3 Design Rules

Three rules follow. The first is to specify, for every intervention, the change that it makes to native capacity, and to track the product of the factors of Equation (20.1) over the sequence, so that the cumulative effect is visible. The second is to set a floor on the grace period and on a for the whole architecture, not for each part. The third is to preserve a mechanism of restoration, so that ρ_{rec} is not zero. Chapter 29 uses these measures to compare the long-run lineages of the body.

20.4 Feasibility, Permission, and the Limits of Morphological Freedom

The ability of one person to withdraw from an intervention is a matter for that person's own capacity and the support available to them. A separate question is whether a society can absorb a class of interventions while keeping its shared institutions usable, and this section addresses it. It also places a limit on how the long-horizon chapters of Part VII are to be read: a possible future body can be worth analyzing even when the case for permitting it is weak.

Definition 20.2 (Four judgments about a bodily form)**.** The ability to produce a bodily form establishes its *technical feasibility* under stated conditions. A case for *present use*

requires evidence that the intervention serves a specified purpose under present conditions. A case for *necessity* requires showing that the purpose cannot adequately be served by available alternatives. A case for *legal permission* must also account for effects on the person, on others, and on the institutions that make the intervention possible.

The four judgments can diverge. A design may be technically feasible and useful in a particular setting and still be unsuitable for general deployment, and the analysis of a design in this book, including every operating range computed in Chapter 30, establishes at most the first of them.

Aesthetic Aversion and Public Justification

Forms such as the extreme engineered body plans imagined in a well-known speculative work of future natural history [58] provoke strong aversion in many readers, and the aversion is a legitimate personal response. It is not by itself a public justification. A public justification for a prohibition must specify what implementing a design would require of the people who inhabit it and of everyone who must recognize, communicate with, care for, or rescue them. The position developed here is that deliberate engineering of humans into forms of that kind should be prohibited, and that a strong presumption should stand against realizing the more extreme body plans of the kind just mentioned, on the grounds set out below and not on grounds of repugnance.

Societal Cognitive Load

Definition 20.3 (Societal cognitive load)**.** The cognitive load that a design places on a shared institution is the additional time, information, and specialized judgment that a responder, clinician, or other participant requires in order to make a consequential decision correctly about a person with that design.

The definition is concerned with what must be known and decided, and not with how the design looks. Consider two designs. The first is visibly unfamiliar, but a responder can read its support requirements through a common interface, so the additional information needed is a single standard lookup. The second looks entirely ordinary, but responding correctly requires knowledge of a proprietary implant that nobody at the scene can identify, so the additional information is unavailable and the required judgment cannot be made in time. The second can impose the greater cognitive load. The comparison separates an aesthetic preference from a burden that legislation could assess.

Cumulative Deployment

There is a precise connection to the mathematics of the book. Suppose an emergency service can maintain C hours of specialized training per year, and two designs require separate protocols that take c_1 and c_2 hours. Then $c_1 \le C$ and $c_2 \le C$ do not imply $c_1 + c_2 \le C$. With $C = 40$ and $c_1 = c_2 = 30$, each design passes a regulatory test in isolation and the pair exceeds the shared capacity by 20 hours. With n designs of equal cost c requiring separate protocols the service can support at most $\lfloor C/c \rfloor$ of them, which is 6 for $c = 6$ hours. This proves a narrow point about cumulative deployment under the stated assumptions. If a

common protocol covers both designs, the cost model changes: with a common protocol of 35 hours the pair fits within C. The legislative question is therefore how the proposed forms interact with the actual institutions expected to support them.

The Regulatory Limit

Where bodily diversity creates additional demands, a society can often meet them through common interfaces, revised protocols, training, and accessible infrastructure. A proposed intervention crosses the regulatory limit when it creates a serious, foreseeable failure of essential services that remains after feasible accommodations have been considered. For irreversible interventions imposed during development, the inquiry also includes whether the resulting person will have practical ways to revise their dependence, change providers, and participate in ordinary civic life, which connects this section to the grace period and independent maintenance of the preceding sections.

The cognitive-load argument can justify prohibiting particular interventions under specified conditions. A categorical ban on a whole family of forms would need an additional principle, for example a prohibition on intentionally imposing irreversible developmental changes that substantially foreclose the resulting person's future agency. Such a principle would also place familiar-looking designs under scrutiny, including some of the developmental proposals of Chapter 19, and a book that adopts it must apply it to its own proposals. The argument of this section stops short of that principle, and the choice of whether to adopt it is a normative one that the argument does not settle.

Evidence Status

Established physiology. That disuse reduces the capacity of tissue, as in bone and muscle, is established [15]. Demonstrated engineering. Dependence on support is documented for dialysis, for ventricular assist devices, and for parenteral nutrition [22, 17]. Proposed design. The model of Equation (20.1), the three measures of the preceding section, and their application to the proposals of this book are developed in the manuscript on continuous microsurgery [9].

Exercise 20.1. For $c_0 = 2$ and a fixed $\delta = 0.07$, find the number of interventions after which the person is no longer self-sufficient, and the grace period after ten interventions for D_{lim} equal to three months of demand.

Exercise 20.2. Suppose that native capacity can be restored at a rate of 0.02 per month by rehabilitation. Starting from c_{10} of Worked Example 20.1, determine whether the person can return to $c \geq 1$ before the damage accumulated during a withdrawal of support reaches D_{lim}, and find the longest withdrawal that permits it.

Exercise 20.3. Propose a way to represent the loss of independent maintenance availability a along a sequence of interventions, and use it to compare an implanted design with a worn one.

Part VI

Aging, Repair, and the Recovery of Admissible Futures

Chapter 21

Computational Physiology and the Simulator

A specification of the kind introduced in Chapter 4 can be checked against a model long before any device exists. Whole-body physiological simulation is the tool for doing so, and the book relies on it for the exercises that ask students to trace budgets over an ordinary day, a period of exertion, and an interruption. This chapter describes the role of a simulator in the design process and its limits, with reference to Physiome, a whole-body physiology simulator written in the Rust language and cited in the manuscript A Relativistic Theory of Longevity as a way to explore physiological regulation and repair computationally [41, 45].

21.1 What a Simulator Is For

Integrative models of human physiology have a long history. The circulatory regulation model of Guyton and colleagues treated the cardiovascular system as a set of interacting control loops, with the kidney’s handling of sodium and water setting long-term arterial pressure [48], and later environments such as HumMod assembled thousands of variables into a model that can be perturbed and observed [49]. A simulator in this tradition serves three purposes in the design of artificial organs. It checks that a proposed device keeps the state inside the viable set $\mathcal{K}$ under the disturbances $\mathcal{D}$ of a specification. It exposes interactions between loops that are not visible when each loop is analyzed separately, as in the competition between water retention and heat rejection of Chapter 7. And it supplies a harness in which perturbations can be applied systematically, which is the use required by Chapter 23.

A simulator organized for this purpose has four parts. The state vector holds the compartments of water and solutes, the thermal state, gas content, and the quantities that describe repair capacity. Physiological modules update these according to conservation laws and control relations. Device modules are attached at the boundaries of Chapter 3, each declaring the streams it exchanges with the body. A perturbation harness applies a scripted disturbance and records the response. The separation between physiological and device modules allows the same body to be tested with different architectures, which is the comparison that the compaction problem of Chapter 17 requires.

21.2 Verification and Validation

Two checks have to be kept distinct. Verification asks whether the code solves the model equations correctly, and it can be done by comparison with analytic solutions. Validation asks whether the model equations describe the body, and it can only be done against measurements.

Worked Example 21.1 (Verifying the garment model)**.** The sweat-recovering garment of Worked Example 7.1 has an analytic prediction: with a heat load of 400 W, a surface rejection of 144 W at an air temperature of 35 °C, and a body heat capacity of 2.43×10^5 $\mathrm{J\,K^{-1}}$, the core temperature rises at $256/2.43 \times 10^5 \approx 1.05 \times 10^{-3}$ $\mathrm{K\,s^{-1}}$ and reaches 2 K in about 1900 s. A forward Euler integration with a time step of one second returns 1898 s, which verifies the integrator for this case to within a tenth of one percent.

A second test checks the control behavior. For bare skin at 35 °C with a proportional sweating response of 300 $\mathrm{W\,K^{-1}}$ and a ceiling of 1.5 $\mathrm{L\,h^{-1}}$, the model settles at the temperature excess for which the evaporative loss balances the load of 414 W, which is $414/300 \approx 1.4$ K, and a four-hour integration gives 1.38 K. The ceiling of 1.5 $\mathrm{L\,h^{-1}}$ corresponds to about 1010 W and is not reached, so the saturating branch is not exercised, and a further test with a larger load would be needed to cover it. These are tests of the code. They say nothing about whether a gain of 300 $\mathrm{W\,K^{-1}}$ describes any person, which is a validation question.

21.3 Limits of a Simulator in the Design of Organs

A simulator can show that a design fails, and it can show that a design survives the disturbances that its authors thought to include. It cannot show that a design improves the life of a person, because the model contains the assumptions of its authors and omits the processes that were not anticipated. The standard of Chapter 4, that a specification states what would have to be measured, therefore applies to the simulator as it does to the device: each result should be accompanied by the measurement that would test it in a body. Code that is published openly, with its parameters and its tests, can be checked by others, and this is a reason to favor simple, well-documented models over elaborate ones.

Evidence Status

Established physiology. Integrative physiological modeling is an established field [48, 49]. Demonstrated engineering. Whole-body simulators exist and are used in research and teaching. Proposed design. Physiome and its use for the design questions of this book are described in the cited work [41, 45], and no clinical validation is claimed here.

Exercise 21.1. Specify the state vector and the device-module interface for a simulator that can reproduce Worked Example 7.1. List the parameters that would have to be validated against measurements on people, and the measurement for each.

Exercise 21.2. Design a verification test for the saturating branch of the sweating response in Worked Example 21.1, giving the load, the expected steady state or the expected rate of rise, and the tolerance.

Chapter 22

Compensation and the Loss of Regulatory Loops

An organism can look functional while it loses some of the loops that sustain it. When one of several redundant regulators fails, the others take up the load, the regulated variable stays where it was, and no familiar function visibly fails. A body in this state has lost something, because its reserve is smaller and its response to a disturbance is slower. Theoretical work on the architecture of autonomy, on topological descriptions of therapy, and on closure among regulatory cycles addresses this situation [38, 46, 39], and this chapter gives a minimal model of it.

22.1 Redundant Loops and Reserve

Let a variable be regulated by n equivalent loops, each able to supply a corrective effort up to c, and let the demand on the regulation be $d(t)$, with resting demand d_r and peak demand $d_{\max}$. While $n_w c \geq d(t)$, where n_w is the number of working loops, the variable is held at its set point. Define the reserve ratio of Equation (4.1) as $\rho = n_w c / d_{\max}$ and the spare capacity at rest as $\sigma = n_w c - d_r$. After a perturbation, the rate at which the variable returns to its set point is proportional to the spare capacity, so the recovery time is

$$\tau_{\mathrm{rec}} \propto \frac{1}{n_w c - d_r}. \tag{22.1}$$

The model is a caricature, and it has the property that the quantities that are visible, namely the value of the regulated variable at rest, do not change until the spare capacity is gone, while the quantities that are not visible, the reserve and the recovery time, change at once.

Worked Example 22.1 (Five loops)**.** Take $n = 5$, $c = 0.6$, $d_r = 1$, and $d_{\max} = 2.5$ in units of effort. With all loops working, capacity is 3.0, the reserve ratio is 1.2, and the spare capacity at rest is 2.0. With one loop lost, capacity is 2.4, which is below $d_{\max}$, so the person can no longer meet the peak demand, while the regulated variable at rest is unchanged. With two lost, capacity is 1.8 and the spare capacity at rest is 0.8. With three lost, it is 1.2 and the spare capacity is 0.2. With four lost, capacity is 0.6, which is below the resting demand, and the variable drifts. Recovery times from Equation (22.1), relative to the intact case, are 1, 1.4, 2.5, and 10 for zero, one, two, and three loops lost, taking the proportionality constant fixed.

The resting value of the variable is normal through three losses out of five. The person loses peak capacity at the first loss and takes ten times as long to recover from a perturbation at the third. The first visible failure at rest occurs only after four of five loops are gone, by which time the loss is nearly total.

22.2 Why Compensation Hides Decline

Compensation hides decline because it is what the system was built to do. The effect is general, and it is the reason that screening by resting values can mislead. It is related to the observation that systems approaching a critical transition recover more slowly from small perturbations, so that the rate of recovery can act as an early warning [55], and to the study of robustness in biological systems, in which loss of redundancy is distinguished from loss of function [51].

For artificial organs the consequence is twofold. A device that replaces a loop can mask a decline in the remaining loops by holding the variable in place, and a design that is evaluated only by the resting value of the variable will appear to succeed. And an intervention that restores a resting value says nothing about whether the lost loops have been restored. The measurement that distinguishes the cases is the response to a perturbation, which is the subject of Chapter 23.

Evidence Status

Established physiology. Compensation among redundant regulators and the slowing of recovery near a transition are documented in physiological and ecological systems [55, 51]. Demonstrated engineering. No artificial organ has been evaluated by the five-loop model. Proposed design. The theoretical treatments of autonomy, topology of therapy, and closure [38, 46, 39] are proposals, and the model of this chapter is a teaching device.

Exercise 22.1. Extend Worked Example 22.1 to loops of unequal capacity, and find the order of loss that conceals the decline longest at rest and the order that reveals it soonest.

Exercise 22.2. A device holds a regulated variable at its set point by supplying a constant corrective effort. Show using Equation (22.1) how the device changes the apparent reserve and the recovery time, and state what test would distinguish a person with intact loops from one whose loops are masked by the device.

Chapter 23

Rejuvenation as Recovery of Admissible Futures

A proposal for what should count as rejuvenation distinguishes three things that are often run together [43]. The first is the reversal of an aging marker, a measured quantity that changes with age. The second is the restoration of a broader biological state, in which a set of such quantities moves toward its earlier values. The third is the recovery of what the organism can actually do under stress, which is a statement about its admissible futures: the set of states it can reach, and the disturbances from which it can return. The proposal is to ask whether an intervention restores responses that had been lost, by applying controlled perturbations.

23.1 Admissible Futures

For a person in state x, with dynamics under control and disturbance as in Appendix D, let $V(x)$ denote the set of disturbances from which the state returns to a nominal region $\mathcal{K}_0$ within a stated time $\tau_{\max}$. This is the set of futures that remain admissible, in the sense that the person can still recover. A perturbation battery $P = \{p_1, \dots, p_m\}$ is a finite set of graded, reversible challenges drawn from ordinary life, for instance a heat load, a fluid deficit, a postural change, a graded exertion, a glucose load, and a cold exposure. The recovery index is

$$\rho_R(x; P) = \frac{1}{m} \sum_{j=1}^{m} w_j \, \mathbf{1}[\tau_j \le \tau_{\max,j}] \,, \tag{23.1}$$

where τ_j is the observed recovery time from perturbation p_j, $\tau_{\max,j}$ is its tolerance, and the weights w_j sum to m and express the importance of each challenge to the specification. The distribution of the τ_j is a finer measure than the index, because the shift from fast to slow recovery occurs before the index changes.

23.2 Three Levels of Claim

The three levels can dissociate. An intervention may change a marker without changing the recovery index, change several markers together without improving recovery, or improve recovery with little change in markers. A claim at the third level can be made only on the basis of the third kind of measurement.

Worked Example 23.1 (A battery before and after)**.** A battery of six perturbations is applied before and after two interventions, and the outcomes below are constructed to show how the levels can dissociate and are not measurements. Before either intervention the person recovers within tolerance from three of the six, so $\rho_R = 0.5$ with equal weights. After

intervention A, which targets a marker, the marker is restored to its earlier value, and the person recovers from the same three, so $\rho_R = 0.5$: the first level of claim is met and the third is not. After intervention B, which restores a missing loop of the kind discussed in Chapter 22, the person recovers from five of the six, so $\rho_R \approx 0.83$, and the marker is unchanged. The first intervention would be judged a success by the marker and the second a failure, and the judgment is reversed if the standard is recovery.

A count is not enough, however, because it can rise while a capacity is lost. Let intervention C restore recovery from three perturbations that were failed before and abolish it for two that were passed, giving $\rho_R = (3 - 2 + 3)/6 \approx 0.67$ against 0.5. The index rose and yet the person can no longer recover from two challenges they could handle. The comparison therefore uses sets and not counts. Let the success set $S(x)$ be the set of battery perturbations p_j for which the state stays in $\mathcal{K}$, returns to $\mathcal{K}_0$ by the deadline $\tau_{\max,j}$, and remains there for a stated dwell time. An intervention is said to expand the admissible futures only when the post-intervention success set contains the baseline set and adds at least one perturbation, $S_{\text{pre}} \subsetneq S_{\text{post}}$. Intervention B meets this test if it retains the three passed perturbations, and intervention C does not, whatever the index says.

23.3 The Test for an Artificial Organ

The proposal supplies a direct test for an artificial organ: after the intervention, what can the person recover from that they could no longer recover from before? An organ that restores a resting value and leaves recovery unchanged has supplemented an operation and has not restored a capacity. An organ that makes recovery possible from perturbations that were previously unrecoverable has changed the set of admissible futures. The test applies to a device as well as to a biological intervention, and it provides the measurement that the exercises of earlier chapters ask students to specify.

Three precautions apply. The perturbations must be graded and reversible, with stopping rules, because a test of recovery places the person under load. The battery should include perturbations that the intervention was not designed to address, since those are where compensation is most likely to be exposed. And the baseline should be measured more than once, because recovery times vary between occasions.

Evidence Status

Established physiology. That recovery rate from perturbation changes with the state of a system and that markers can dissociate from function are established concepts [55]. Demonstrated engineering. Challenge tests of physiological function, such as glucose tolerance and graded exercise tests, are standard clinical tools. Proposed design. The recovery index of Equation (23.1), the framing in terms of admissible futures, and the use of the battery as the test of rejuvenation are proposals [43], which appear as a section of the cited manuscript.

Exercise 23.1. Design a perturbation battery for the sweat-recovering garment of Chapter 7. State each perturbation, its grading, its stopping rule, and its recovery tolerance.

Exercise 23.2. Show that the recovery index of Equation (23.1) cannot detect a shift from recovery in one minute to recovery in ten minutes if both are within tolerance. Propose a modification of the index that can.

Chapter 24

Turnover, Damage, and Retained Mechanical History

Tissues differ greatly in how quickly their components are replaced, and this governs two questions of the book: how long the effects of earlier loading are retained, and how damage accumulates when repair is limited. The first is the subject of the manuscript on Mechanical Inheritance, which asks how long tissue retains the effects of earlier mechanical loading and how turnover and accumulated changes affect later function [12]. The second is treated in a manuscript on recursive continuation, which uses cellular senescence as a case in which damage accumulates through ordinary activity while the resources for repair are constrained [42].

24.1 Retention and Turnover

If a component is replaced at random with a half-life $t_{1/2}$, the fraction of the original material remaining after a time t is

$$f(t) = 2^{-t/t_{1/2}}. \tag{24.1}$$

Measurements of the age-related accumulation of modified collagen give very different half-lives for different tissues, about 15 years for skin collagen and more than a hundred years for the collagen of articular cartilage [56].

Worked Example 24.1 (Retained traces at age forty)**.** With the half-lives $t_{1/2} \approx 14.8$ and 117 years, Equation (24.1) gives, at age forty, $f = 2^{-40/14.8} \approx 0.15$ for skin and $f = 2^{-40/117} \approx 0.79$ for cartilage. Here f is the fraction of collagen present at birth that remains at forty. It is not the fraction of the collagen present at forty that was laid down in the first four decades: with constant deposition and exponential loss that fraction is $(1 - e^{-\lambda t})/(\lambda t)$, which is about 0.45 for skin and 0.89 for cartilage. The cohort survival of 0.15 and 0.79 is a separate quantity, and the hypothesis that a developmental signature is retained in proportion to it is an additional assumption. The prediction that follows is that slower-turnover tissues should retain stronger traces of developmental loading into adulthood. The assumptions are strong: replacement is not uniform, some pools of collagen are much longer lived than the average, and modification accumulates in addition to retention, so the numbers fix an order of magnitude and are not a forecast for any tissue.

The consequence for a developmental integument, as in Chapter 19, is that the tissues most strongly shaped by early mechanical conditions are those with the slowest turnover, and any change made early will be retained longest in them. This is both an advantage, for conditions that should persist, and a hazard, for conditions that prove unsuitable.

24.2 Damage Under Constrained Repair

Let the damaged population N be produced at rate s by ordinary activity and cleared at a rate that saturates, so that the clearance resource is shared among the damaged items,

$$\frac{\mathrm{d}N}{\mathrm{d}t} = s - \frac{\kappa_0 N}{1 + N/N_0}. \tag{24.2}$$

Clearance saturates at $\kappa_0 N_0$. A steady state exists only when $s < s_c \equiv \kappa_0 N_0$, and then

$$N^* = N_0 \, \frac{s/s_c}{1 - s/s_c}. \tag{24.3}$$

Proposition 24.1 (Threshold for damage accumulation)**.** Let $N(0) = N_{\mathrm{init}} \geq 0$ in Equation (24.2), with $s, \kappa_0, N_0 > 0$. If $s < s_c$ then $N(t) \to N^*$ monotonically. If $s > s_c$ then $N(t) \to \infty$. If $s = s_c$ then $N(t) \to \infty$ as well, but only like $\sqrt{2 N_0 s_c t}$.

Proof. The right-hand side $g(N) = s - \kappa_0 N/(1 + N/N_0)$ is continuous, with $g(0) = s > 0$ and g strictly decreasing toward $s - s_c$ as $N \to \infty$. If $s < s_c$ it has a single root N^* with $g > 0$ below and $g < 0$ above, so solutions are monotone and converge to it. If $s > s_c$ then $g \geq s - s_c > 0$ for all N, so $N(t) \geq N_{\mathrm{init}} + (s - s_c)t$. If $s = s_c$ then $g(N) = s_c/(1 + N/N_0) > 0$, and N increases without bound since $g > 0$ everywhere; for large N, $g \approx s_c N_0/N$, so $N \, \mathrm{d}N \approx s_c N_0 \, \mathrm{d}t$ and $N^2 \approx 2 N_0 s_c t$. □

Below the critical production rate the burden is bounded, and at $s/s_c = 0.5$ it equals N_0, at 0.9 it equals $9N_0$, and at 0.99 it equals $99N_0$. Near the critical rate a small increase in production, or a small decrease in the clearance capacity, produces a very large change in the burden, and beyond it the burden grows without bound. The model is a version of the idea that the accumulation of damage reflects a competition between production and a limited repair resource [50], with cellular senescence as a case in which damaged cells persist and alter their surroundings [37].

The relevance for artificial organs is that a device or interface which raises the production of damage, for example through chronic mechanical irritation or inflammation at a margin, or which draws on the same repair resource, moves the system toward s_c. The nonlinearity of Equation (24.3) means that the effect cannot be extrapolated from small trials: a modest effect on production may be harmless far from the threshold and severe near it. The erosion model of Chapter 20 is a companion, in which capacity declines in steps, and the two models together describe a person whose repair resource is both shrinking and more heavily used.

Evidence Status

Established physiology. Turnover rates of collagen in different tissues and the accumulation of senescent cells with age are documented [56, 37], and the competition between repair and damage is a standard theme [50]. Demonstrated engineering. Measurements of tissue turnover in humans exist, and measurements of the retention of developmental loading do not. Proposed design. The prediction of stronger retention in slower-turnover tissues and the use of the senescence model for devices are proposals [12, 42].

Exercise 24.1. Compute f from Equation (24.1) for a tissue with half-life 5 years and one with 40 years at ages 10, 40, and 70. Design a measurement in an animal model that would test the prediction of stronger retention in the slower tissue.

Exercise 24.2. For Equation (24.2), find the factor by which the steady-state burden N^{*} rises when the production rate rises by ten percent from $s/s_c = 0.8$ and from $s/s_c = 0.95$. State what this implies for the design of a safety margin.

Chapter 25

Maintenance Without Memory: Cohort Turnover in Complex Systems

The previous chapter treated the turnover of molecules and cells. A system that is maintained over decades also turns over its maintainers, and what a maintainer knows about why something was built is retained in the same way that a modified collagen is retained: it persists in proportion to the survival of the carriers. A habitat-carrying body that is meant to last a century will be serviced by a succession of people, none of whom made the original design choices. This chapter develops a short model of that succession and uses it to examine claims made by the programmer and game developer Jonathan Blow in a recent interview about why software grows more complicated and less reliable [57]. The claims are the testimony of one experienced practitioner and not measurements, and they are used here as hypotheses that a simple model can make precise.

25.1 Experience in a Growing Population

Let the number of people maintaining a system be $N(t) = N_0 e^{gt}$, and let each person leave at rate μ per year. New people enter at rate $(g + \mu)N(t)$, so that the population grows at g while replacing leavers. The tenure of a person drawn at random from those present at time t then has the exponential density

$$p(a) = (g + \mu)\, e^{-(g+\mu)a}, \tag{25.1}$$

with mean $1/(g + \mu)$, median $\ln 2/(g + \mu)$, and a fraction $e^{-(g+\mu)\tau}$ having at least τ years of experience.

Worked Example 25.1 (Fast growth and the missing senior cohort)**.** In the interview a field that doubles every year is used to show that most of its members are inexperienced [57]. In the continuous form of the model, doubling means $g = \ln 2 \approx 0.69$ per year, and with $\mu = 0$ the mean tenure is $1/g \approx 1.4$ years, the median is exactly one year, and the fraction with at least five years is $e^{-5g} = 2^{-5} \approx 3\%$. A field growing at ten percent a year, with the same $\mu = 0$, has a mean tenure of ten years and a fraction of $e^{-0.5} \approx 61\%$ with at least five. A stable population with $g = 0$ and $\mu = 0.1$ has the same mean tenure of ten years, and $e^{-1} \approx 37\%$ have at least ten. The same arithmetic applies to any maintenance workforce: it is the sum of the growth rate and the departure rate that sets how much experience is on hand, and either one alone can empty the senior cohort.

The practical consequence concerns what counts as normal. The person with the median tenure has seen only the system as it is, and a practice that has been present for the whole of that tenure appears to be a feature of the world, not a decision that someone made. When the median tenure is a year or two, the reference point against which new proposals

are judged moves forward at nearly the speed of the system itself, and the reasons that earlier choices were made are known to a small fraction of those present.

25.2 The Half-Life of a Rationale

Consider a design decision made at time s, whose rationale is held by the people present at s and is not written down. At a later time t the number of those people still present is $N(s)e^{-\mu(t-s)}$, and the population is $N(s)e^{g(t-s)}$, so the fraction of the current maintainers who were present at the decision is

$$\phi(t-s) = e^{-(g+\mu)(t-s)}, \tag{25.2}$$

which is the tail of Equation (25.1). This is the retention law of Equation (24.1) with a half-life $\ln 2/(g+\mu)$. For a stable team with $\mu = 0.2$ per year the half-life of an undocumented rationale is $\ln 2/0.2 \approx 3.5$ years, and a decision made thirty years ago is known to $e^{-6} \approx 0.2\%$ of those present, while a growing team loses the rationale faster. Documentation changes the carrier. It makes the rationale a property of the record, which is retained for as long as the record is kept and read, and not of the people, but a record is read only when someone thinks to ask, and the person who does not know that a decision was made does not know to look for it.

Two features of the interview fit this model. The first is its account of the early web, which is described as not having been intended to become what it is, with the original discussions unknown to most of those now working on it [57]. The second is the remark that a person with five or six years of experience is among the more experienced people around and still does not know much of what lies below, so that the system resembles a landfill whose deeper layers are inaccessible. In terms of the model, the loss is not a fault of the newer people. It is the expected result of the rates g and μ.

25.3 Hardening Under Accumulation

The interview also makes a claim about the changeability of software: it becomes harder to change as it grows, because the larger it is, the more likely it is that a change has to touch other parts [57]. A simple model makes this concrete. Components are added one at a time, and each new component depends on each earlier one independently with probability p. A component added at rank j among n is then depended on by about $p\,(n-j)$ later components, so the expected number of dependents of a component chosen at random is about $pn/2$, and the expected cost of changing it, counted as the number of components that have to be examined, is of order $1 + pn/2$. In a branching approximation, in which each affected component spreads the change to its own dependents, a change to an early component can cascade to a large part of the system once the mean number of dependents of the components involved exceeds one, which here occurs when pn is of order one or larger. The cost of change therefore grows with size even when each addition is made carefully, and a system that keeps the ratio of dependency to size constant, by limiting p as n grows, is the only kind that does not harden. This is the mechanism by which the loss of the rationale becomes consequential: the people who would be able to say which dependencies are real are those who are no longer present.

25.4 Complexity as a By-Product of Incompatibility and of Standards

The interview gives two further reasons that systems become more complicated than anyone wants, and both can be stated as models.

The first is the incompatibility of platforms. Its account is that the vendors of large platforms have an interest in making the interfaces for their platform differ from those of others, so that software written for one is costly to move, and the practice is now universal among them [57]. If a developer who has used k platform-specific interfaces faces an exit cost $c\,k$, and the vendor gains from the exit cost, then every new interface that is a little different from the corresponding one elsewhere is rewarded. Nobody in this account needs to want complexity. It is produced by the incentive to raise the cost of leaving, and the interview notes that the same complexity growth occurs in a system, Linux, to which this incentive is said not to apply, so that the incentive cannot be the only cause.

The second is the structure of standards. A standard must be usable by everyone who adopts it, so its content is the union of what each adopter needs, while each adopter uses a small part. Suppose that each of k adopters needs a random subset of s features from a universe of M. The expected size of the union is $M\big[1-(1-s/M)^{k}\big]$. For $M = 100$, $s = 10$, and $k = 20$ the union has about 88 features, and a typical adopter uses 10 of them, so about 11% of what each must implement and understand is used by them, and the rest exists for the sake of others. The interview describes this effect with the example of the multiple runlevels of a Unix initialization standard, and it traces the later growth of the boot and service manager to a size of over a million lines, a figure that it gives as recalled and that is not verified here [57].

25.5 Games as a Contrast Case

The interview treats video games as the most complex software that its speaker knows of, and says that games work against the received paradigms of software engineering, so that the more fun a game is the more it departs from them [57]. This book draws a narrower lesson from the contrast. A game is exercised constantly by people who notice immediately when it fails and who will leave it if it does not do what they want, so its complexity is subject to a test that is applied on every use. The complexity of an operating-system layer is tested only when something breaks, and often by someone who is unable to trace the cause. The distinction does not concern the amount of complexity, and what matters is whether the complexity is under a continuous correction signal.

25.6 Consequences for a Body That Carries Its Habitat

The three mechanisms bear on the architectures of Chapters 13, 14, and 30 in the following way. A body with an internal auxiliary physiology and an exchange interface is a system of components with dependencies, maintained by a succession of clinicians and technicians. For a service life of a century and a clinical workforce with $\mu = 0.1$ per year, the fraction of present maintainers who were present at the design is $e^{-10} \approx 4.5 \times 10^{-5}$, and the rationale

of any undocumented choice is gone within a few decades. The grace period of Chapter 20 measures how long the body can be without a physical support, and it has an institutional counterpart, the interval for which the maintenance of the device can be carried on without the original knowledge. Four design proposals follow, and they are proposals and not results.

The first is that the number of interfaces between the body and its maintainers should be small and fixed by function. The Umbilicus of Chapter 14 is such an interface, and a design that gave every device its own connector would raise p in the hardening model. The second is that each device should carry its rationale, in the sense of a record stored with the device and readable without the original makers, so that the carrier of the rationale is not the workforce. The third is that the standard for the interface should be minimal, with extensions kept outside the core, to limit the union effect. The fourth is that the correction signal should be built in: a single person is a user base of one, so the test of Chapter 23, the recovery battery, plays the part that continual play plays for a game, and it should be applied routinely and not only after a failure.

Evidence Status

Established physiology. The relations among growth rate, departure rate, tenure, and the survival of undocumented knowledge are standard results for an exponentially growing, memoryless population. Demonstrated engineering. The statements about software are the observations of one practitioner in an interview and have not been measured here, and the figures of Worked Example 25.1 are the outputs of the stated models. Proposed design. The application to the maintenance of an artificial organ over a century, and the four design proposals, are proposals developed in this book.

Exercise 25.1. For a maintenance workforce with $g = 0.05$ and $\mu = 0.15$ per year, find the median tenure and the half-life of an undocumented rationale. Then find the departure rate at which the rationale half-life of a stable team equals the service life of the longest-lived component of a device, taken as fifteen years.

Exercise 25.2. In the hardening model, let the dependency probability be $p = c/n$ for a constant c, so that each component has a bounded expected number of dependents. Show that the expected number of components touched by a change no longer grows with n, and explain what design discipline this represents.

Exercise 25.3. Using the union formula for standards, find the number of adopters k at which the typical adopter uses less than five percent of the standard when $M = 100$ and $s = 10$. State what a modular standard with a minimal core and optional extensions does to the ratio.

Chapter 26

Restoration Before Explanation

A function can be restored before its mechanism is understood. The history of artificial and transplanted organs contains many cases in which a procedure worked, in the sense that the person lived, before the processes that made it work were characterized, and some cases in which the explanation came later and changed the practice. A manuscript on restoration before explanation takes this as its starting point and relates kidney transplantation, organ-system restoration, and bioelectric morphogenesis to the idea that restoration of function can precede its explanation [44]. This chapter draws out the consequences for how claims about artificial organs should be graded.

26.1 Examples

The first successful kidney transplant between identical twins in 1954 restored renal function in a person who would otherwise have died, in a setting in which the immunology of rejection was still incompletely understood [53]. Dialysis was in use before the dynamics of solute removal were fully described in the terms used in clinical practice today. In the study of development and regeneration, the manipulation of bioelectric states has been reported to redirect the formation of structures in animal models, with the mechanism by which the electrical pattern is read still under investigation [52]. In each case the restoration is real and the explanation is partial.

26.2 Three Tiers of Evidence

The cases suggest three tiers, which a claim about an artificial organ should be graded against. At the first tier, function has been restored under a perturbation battery of the kind described in Chapter 23. At the second tier, the mechanism by which the device acts has been characterized, so that its action can be described without reference to the outcome. At the third tier, the mechanism is sufficient to predict the failures of the device, including those not yet observed. A device at the first tier alone has a record and no theory, and its failure modes are known only from the failures that have occurred.

26.3 Design Consequences

A device at the first tier alone requires more conservative provisions, because its failures cannot be anticipated. The design responses are the ones the book has developed: a longer grace period, a retained fallback, an ability to leave in the sense of Chapter 20, and a monitoring scheme that detects slowing of recovery before function fails. The converse also holds. A device with a well-characterized mechanism that has not been tested under a

Table 26.1: Tiers of evidence and the claims they support.

Tier	What has been shown	What may be claimed
Restoration	Function recovers under a perturbation battery	That the device works for the conditions tested
Mechanism	The action of the device is characterized	That the effect can be reproduced and that variants can be designed
Prediction	The mechanism predicts unobserved failures	That safety margins can be set in advance

battery has a theory and no record, and its claims should be graded accordingly. Neither tier substitutes for the other.

Evidence Status

Established physiology. The history of kidney transplantation and dialysis is documented [53]. Demonstrated engineering. Bioelectric control of patterning has been reported in animal models [52]. Proposed design. The use of the three tiers to grade claims about artificial organs, and the manuscript from which the framing is taken [44], are proposals.

Exercise 26.1. Assign each of the devices discussed in Chapter 2 to a tier in Table 26.1, and state what additional evidence would raise it to the next tier.

Exercise 26.2. Argue for or against the claim that a device at the restoration tier only should be accepted for use by a person with no alternative. State what conditions on monitoring and reversibility you would require.

Chapter 27

Further Proposals in Architecture, Repair, and Cognition

A collection of more speculative lines, gathered in archived notes on life extension, bears on the architectures of this book [40]. Three are taken up here, each with a statement of what a specification would need and how the tools of the book would be applied. The notes describe the lines at the level of ideas, and no specification or measurement is available for any of them.

27.1 Forest-Inspired Organ Architecture

Forests and trees are distribution networks that move water over great distances against gravity with no central pump, using branching conduits and gradients in water potential, and the idea of taking forests as a model for the architecture of organs connects to the author's work on xylomorphic computation [47]. A quantitative handle on branching networks is Murray's law, which states that at a bifurcation of a vessel carrying laminar flow, the cube of the radius of the parent equals the sum of the cubes of the radii of the daughters when the network minimizes the sum of pumping power and the metabolic cost of maintaining its volume [54]. For a symmetric bifurcation this gives a daughter radius $r_1 = r_0/2^{1/3} \approx 0.794\, r_0$.

The law gives a way to check a proposed distribution mesh for the fluid recovery of Chapter 13: if the conduits of the mesh depart strongly from the cubic rule, then either the network is not minimizing the cost it is assumed to minimize, or the flow is not laminar and the rule does not apply. Whether the structures of a forest transfer to the interstitial space of a human body, which has no pressure gradient of the kind that tall trees exploit, is a question that the specification would have to answer.

27.2 Tissue Repair as a Traversing Process

A proposal for tissue repair modeled on a process, called Yarncrawler in the author's earlier work, imagines a repair process that moves along a structure and mends it as it goes. Without describing details that are not in the notes, a general constraint applies to any repair process that serves a structure of total length L. If damage arises at a rate μ per unit length per unit time, the total arrival rate of damage events is $\lambda_d = \mu L$. If each repair takes a time t_r and there are N repairing agents, each busy a fraction $\lambda_d t_r / N$ of the time, the agents keep up only if

$$N \geq \lambda_d \, t_r = \mu L \, t_r. \tag{27.1}$$

Below this bound the unrepaired damage grows without limit, which is the same threshold behavior as in Equation (24.2). The bound is a necessary condition. Travel time between sites and the requirement that agents find damage add to it.

27.3 Cognitive Anti-Senescence

The notes also collect proposals directed at the decline of cognitive function with age. The tools of the book require a specification before any such proposal can be assessed: which cognitive functions are to be preserved, under which perturbations, such as fatigue, interruption, or novel demands, and with what recovery tolerance. The compensation model of Chapter 22 predicts that a person can show preserved performance on routine tasks while losing the reserve that supports unusual ones, so the battery should include perturbations that tax reserve. Chapter 12 separated the criteria of surface area, wiring distance, transport, and coordination, and an intervention directed at cognition should state which of them it acts on.

Evidence Status

Established physiology. Murray's law is a standard result in vascular biology [54]. Demonstrated engineering. Branching distribution networks are used in engineering, and cell-based and tissue-engineering approaches to repair exist. Proposed design. The forest-inspired organ architecture, the traversing repair process, and the cognitive proposals are speculative lines recorded in archived notes [40, 47], and the analyses here set conditions on them without establishing them.

Exercise 27.1. Derive Murray's law by minimizing the sum of the pumping power $Q^2 R_h$, with $R_h \propto L/r^4$ the hydraulic resistance of a laminar vessel, and a metabolic cost proportional to the vessel volume $\pi r^2 L$, at fixed flow Q. Show that the optimal flow is proportional to r^3.

Exercise 27.2. For a structure of 0.4 m with a damage rate of 100 events per meter per hour and a repair time of 20 minutes, find the minimum number of agents from Equation (27.1). State what the bound does not include.

Part VII

A Hundred to a Thousand Years

Chapter 28

The First Century

The design horizon of one hundred years is used here as a discipline and not as a forecast. The method is to work backward from services that people might want continuously available, to ask what existing technologies would have to be joined to provide them, and to identify at each joining whether the missing piece is a better component or a different architecture. Each trajectory ends in a stall point, a place where it would stop unless something new were found, and naming the stall points is the main result.

28.1 Three Services

Three services serve as the targets: temperature regulation, mobility and impact support, and water management. Each has an existing form in garments, portable devices, and appliances, and each is a function that the specification of Chapter 4 could require. The question is how these existing forms might gradually share power, sensing, and maintenance, so that the three become one system worn on the body.

Table 28.1: Trajectories toward shared worn services. Each row lists the transition, what it depends on, and where it could stall.

Transition	Starting point	Gated by	Stall point
Heated garment to adaptive thermal layer	Battery-heated jackets; passive weaves	Component and control	Active power for heating exceeds what cells can carry
Padded garment to load-sharing support	Protective padding; powered braces	Component, then architecture	Mass and power of actuators; fit to the moving body
Wicking garment to fluid-managing layer	Wicking fabric; liquid-cooled suits	Architecture	Heat rejection conflicts with water retention
Separate batteries to shared power	Independent batteries in each device	Architecture	Safety certification of a shared bus; standards
Separate maintenance to service renewal	Washing and replacement by the owner	Architecture and institution	No service network that handles worn layers

The table distinguishes two kinds of gating. A component gate is a quantity, such as a specific energy or an actuator efficiency, that has to improve by a stated factor. An architecture gate is a change in how the parts are joined, for which no single component improvement is sufficient. The distinction matters because component gates can be forecast from rates of past improvement and architecture gates cannot.

Worked Example 28.1 (A component gate and an architecture gate)**.** Sustaining 100 W of

electrical heating for eight hours from a 1 kg battery needs a specific energy of 800 $\mathrm{Wh\,kg^{-1}}$. With cells at 200 $\mathrm{Wh\,kg^{-1}}$, as assumed in Worked Example 15.1, the gap is a factor of four, a component gate that does not depend on design choices elsewhere.

An architecture gate is illustrated by pooling. Suppose a cold day requires 240 Wh for heating, 40 Wh for mobility assistance, and 120 Wh for water handling and sensing, and a temperate walking day requires none for heating, 160 Wh for mobility, and the same 120 Wh. If each device carries its own reserve sized for its worst day, the total is $240 + 160 + 120 = 520$ Wh, or 2.6 kg at 200 $\mathrm{Wh\,kg^{-1}}$. A pooled reserve is sized for the worst total day, 400 Wh on the cold day, and weighs 2.0 kg, a saving of about a quarter. The gain depends on the services not peaking together and on the shared reserve being safe and standardized. No cell improvement produces it.

28.2 Stall Points as Results

A trajectory that is traced to its stall point has produced something useful, which is a statement of what would have to be true for the next stage to be reached. The heated-garment path stalls on energy, the fluid-managing path on the coupling between water retention and heat rejection that Chapter 7 analyzed, and the service path on institutions. A design horizon is therefore a map of conditional claims, in which each claim names its condition.

Evidence Status

Established physiology. The energy requirements follow from the physiology of Part II. Demonstrated engineering. Each starting point in Table 28.1 exists as a product or prototype, and the specific energy of present cells is documented. Proposed design. The joined system and the intermediate stages are design proposals, and they are intended as conditional statements and not as predictions.

Exercise 28.1. Add three rows to Table 28.1 for services of your choice, and classify each transition as gated by component, by architecture, or by institution.

Exercise 28.2. Using Worked Example 28.1, compute the pooling saving if the warm exertion day adds 60 Wh for active cooling. State whether the worst total day changes.

Chapter 29

The Next Several Centuries

Over longer spans, a single trajectory is replaced by lineages, families of designs that develop along different lines and answer to different conditions. Three lineages are compared: a wearable lineage that carries its support in worn layers, an implanted lineage that places support in tissue, and an exchange lineage that depends on external infrastructure through interfaces such as the Umbilicus. The chapter introduces the possibility that they coexist, because they answer to different climates, occupations, maintenance systems, and preferences about bodily autonomy.

29.1 Why Lineages Persist Together

Each lineage has a region of conditions in which it does best. The wearable lineage has the advantage of renewal without surgery and of removability, which suits mobile occupations and variable climates. The implanted lineage has the advantage of continuous support with no burden on the surface, which suits settings in which the person cannot attend to equipment. The exchange lineage has the advantage of using large and efficient equipment, which suits settled life with dense infrastructure. Preferences about bodily autonomy cut across these conditions: some persons will accept a permanent opening and others will not accept an implant, and a society in which these preferences are respected will contain all three.

29.2 Disturbances

A lineage is tested by disturbances, and four are used here. Relocation moves the person away from the infrastructure that supports the design. Loss of a provider removes a service, whether a manufacturer, a maintenance network, or a utility. Interrupted power removes the energy on which active parts depend. A change in the materials available for repair alters what can be replaced and from what.

Worked Example 29.1 (Counting single points of failure)**.** A useful comparison is the number of distinct providers that must remain available for each lineage to meet its specification over a year. For the wearable lineage the list includes a supplier of renewal layers and a water treatment point, two providers, and the layers can be replaced from several sources if the design is open. For the implanted lineage, it includes a manufacturer of nodes, a clinical service for implantation and removal, and an energy source if the nodes are not self-powered, three providers, and some of them are unique. For the exchange lineage, it includes a utility, a maintainer for the interface, and a supplier of consumables, again three, but each of them is a facility that serves many persons and is maintained by an institution. The counts are not decisive on their own, and the numbers of alternatives for each provider

Table 29.1: Qualitative response of the three lineages to four disturbances, as indicated by the quantities of Chapter 20.

Disturbance	Wearable	Implanted	Exchange
Relocation	Layers carried; renewal needs a service	Unaffected for function; maintenance needs a specialist	Severe unless infrastructure exists at the destination
Loss of provider	Layers can be self-made if designs are open	Device support ends; native capacity decides the grace period	Severe; grace period set by native capacity
Interrupted power	Passive layers continue; active layers stop	Depends on whether nodes are powered locally	Equipment stops; carried reserves bridge
Change in materials	Textiles are adaptable	Implants need qualified materials	Large equipment is adaptable only with service

are as important as the number of providers, which is the quantity that the threshold model of Chapter 8 treats.

The comparison does not lead to a ranking. It leads to the observation that the lineages differ in the kind of failure they suffer, and that a society which contains all three has a form of redundancy at a higher level than any individual's design.

Evidence Status

Established physiology. The disturbances are generic and the responses are derived from the physiology and the dependence analysis of earlier chapters. Demonstrated engineering. No lineage exists in the mature forms discussed. Proposed design. The lineages and their coexistence are design scenarios.

Exercise 29.1. Add a fourth disturbance, such as a sustained rise in environmental temperature, to Table 29.1 and fill in the entries. Identify which lineage the disturbance favors.

Exercise 29.2. For one lineage, enumerate the providers needed over a year and the number of independent alternatives for each. Use the common-cause relation of Equation (8.4) to estimate how the system reliability is limited by the least replaceable provider.

Chapter 30

The Thousand-Year Body

Over a thousand years, comparison replaces prediction. The comparison establishes at most technical feasibility, and Section 20.4 explains why that is a weaker claim than permission. Three mature architectures are examined under the same accounting exercise for matter, energy, maintenance, and interruption, and then the environment is changed to see which provides the wider range of possible lives. The first architecture is compact and carries much of its habitat in its integument. The second has a distributed internal auxiliary physiology of the hepastitial type. The third depends on a rich external ecology with frequent, reversible coupling.

30.1 The Accounting Exercise

For each architecture the student fills in the four accounts of Chapter 17, the water, heat, and energy ledgers of Chapter 5, the grace periods of Chapter 4, and the measures of Chapter 20. The architectures differ in where the accounts fall. The compact architecture has large worn-hardware and consumables accounts and small exposure. The distributed internal architecture has a large tissue account, a small worn account, and an exposure that is concentrated in the maintenance of the nodes. The external ecology has small carried accounts and large exposure, with a grace period that depends on the frequency of coupling.

30.2 Operating Range

The operating range $R(A)$ of Definition 4.1 is a joint object, and it is contained in the intersection of the ranges of the individual functions. The containment can be strict. Suppose a single actuator, such as a heat-rejection surface, serves two functions, water retention and thermal balance, as in Worked Example 7.1. Each function alone may be satisfied at 35 °C, retention with the surface closed and balance with it open, and yet no single setting of the actuator satisfies both, so the environment lies in the intersection of the individual ranges and outside the joint range. The sweat-recovering garment is such a case. The intersection becomes exact when the functions are separable, meaning that each draws on its own resources so that a policy can serve them independently. Under separability and independent limits in the two dimensions, if function f is satisfied for ambient temperatures in $[a_f, b_f]$ and for a distance from service up to $d_f = v\,\tau_f$, where v is the rate at which the person moves away from service and τ_f the grace period for that function, then

$$R(A) = \Big[\max_f a_f,\ \min_f b_f\Big] \times \Big[0,\ \min_f d_f\Big]. \tag{30.1}$$

The range is limited by the most restrictive function in each dimension, and Equation (30.1) is an upper bound on the joint range in general and an equality only under separability.

The consequence for design is that improving a function that is not the most restrictive does nothing for the range.

For the running problem of Chapter 4, the garment G passes e_1 but fails e_2 under the joint specification of heat and water, the hepastitial subsystem H passes e_1 and e_2 with a reliability reduced by its shared layer, and only U and H with stored reserve can address e_3, which is the three-day interruption, according to their grace periods.

Worked Example 30.1 (Three architectures in two dimensions)**.** For illustration, take the following parameters, which fix the method and are not predictions. The compact architecture has temperature limits $[-25, 38]$ °C and a grace period of 10 days. The distributed internal architecture has limits $[-10, 42]$ °C and a grace period of 3 days. The external ecology has limits $[-15, 35]$ °C and a grace period of half a day, which is the longest uncoupled interval. With $v = 40$ km per day, the distance limits from Equation (30.1) are 400, 120, and 20 km.

No architecture dominates. The compact architecture has the largest span of cold tolerance and distance, the distributed internal architecture the widest tolerance of heat, and the external ecology the narrowest range of conditions, but at the lowest carried mass and with access to the largest equipment while coupled. If the environment is changed so that ambient temperature reaches 40 °C, only the distributed internal architecture remains in range, and if it is changed so that service is available everywhere within 20 km and no one travels beyond, the external ecology matches the others at the lowest carried mass. The statement of which architecture provides the widest range of possible lives has no answer until the set of environments is specified, which is the lesson of the exercise.

30.3 What the Comparison Supports

The comparison supports three conclusions. First, the range of a design is set by its binding constraint, and the binding constraint can be moved only by a change in the corresponding function. Second, the architectures trade carried mass, exposure, and range, and none reaches a corner of all three. Third, a society that supports several architectures preserves a broader set of possible lives than one that supports a single architecture, a claim that depends on the architectures being compatible with one another in shared settings.

Evidence Status

Established physiology. The accounts and ledgers are those of Parts I and II. Demonstrated engineering. None of the mature architectures exists. Proposed design. The parameters of Worked Example 30.1 are illustrative, and the comparison is intended to teach a method.

Exercise 30.1. Change the parameters of Worked Example 30.1 so that the external ecology has the largest range in the distance dimension. State which parameter must change and by what factor, and whether the change is physically plausible.

Exercise 30.2. Add a function with a limit on altitude and recompute Equation (30.1). Identify which architecture's range is reduced.

Chapter 31

Capstone on the Minimum Habitable Body

The capstone asks the student to design a smaller body that preserves an explicitly stated set of functions, and to defend the design with the tools of the book. The work product has four parts: a functional specification, a model, a comparison with existing anatomy, and a century-scale transition path that begins with an electric jacket. The design is evaluated by whether its claims are stated with their conditions, and not by whether it is feasible in the end.

31.1 The Specification

The student states the functions to be preserved, in terms of the viable set $\mathcal{K}$, the disturbance class $\mathcal{D}$, and the duration T of Chapter 4. The specification includes a set of activities, such as walking, carrying a load, and working in a seated posture, a range of environments, a required reserve ratio, a required grace period for each function under interruption of each support, and a requirement for the ability to leave in the sense of Chapter 20. A specification that omits any of these is returned for revision.

31.2 The Model

The model traces the water, heat, and maintenance budgets over three periods: an ordinary day, a period of exertion, and a prolonged service interruption. The ordinary day provides the baseline, and the other two test the margins.

Worked Example 31.1 (A baseline day)**.** A day is composed of eight hours of sleep at a metabolic power of 65 W, fourteen hours of light activity at 85 W, and two hours of exertion at 300 W. The energy dissipated is $8 \times 65 + 14 \times 85 + 2 \times 300 = 520 + 1190 + 600 = 2310$ Wh, or 8.3 MJ, close to the 8.6 MJ of Table 5.1. The exertion period contributes about a quarter of the energy and, in a warm environment, most of the sweat. For the running problem of Chapter 4, with the 300 W of this period and 75 W of external work, the heat to be rejected is 225 W and the required evaporative cooling at an air temperature of 35 °C is $225 + 14 = 239$ W, or 0.35 L h^{-1}, so the exertion period costs 0.7 L of water in sweat. The rest of the day's water follows Table 5.1. The student builds on this baseline and states how the design changes each entry.

The prolonged interruption is modeled by withdrawing each external support in turn and computing the time to exit $\mathcal{K}$, using the grace period relations of the preceding chapters. The result is a table of grace periods by support, and the least entry is the binding constraint on the design.

31.3 Comparison and Transition

The comparison with existing anatomy states, for each function, whether the design performs it at greater, equal, or lesser capacity, and records the four accounts for both. The transition path begins with an electric jacket and proceeds by stages in the manner of Chapter 28, naming at each stage whether the gate is a component, an architecture, or an institution, and where the path could stall.

31.4 Criteria

The work is assessed on six criteria. The specification is complete and measurable. The budgets close, in the sense of Chapter 5, with every stream accounted for. The claims are separated into established, demonstrated, and proposed. The design states where the function belongs: some belong on the body, and some work better at the scale of a room, vehicle, or town, and a successful design can show this explicitly, which is a result and not a failure. The design is graded against the tiers of evidence of Chapter 26. Finally, the student states what would have to be measured to establish that the design improves the life of the person who depends on it, and specifies a perturbation battery in the sense of Chapter 23, so that the improvement is judged by what the person can recover from and not only by resting values.

31.5 What the Book Has Argued

The tools of Part I make an organ a collection of operations with a boundary, and a specification a set of admissible trajectories. The tools of Part II make a design answerable to conservation, control, and scaling. The architectures of Parts III and IV make explicit what each proposal would require. The accounts of Part V show where an apparent saving goes, and the measures of dependence show how a sequence of benefits can reduce options. The horizons of Part VI turn the question of the future into a set of conditional statements. Part VI added the requirement that every claim of improvement be tested by recovery from perturbation and graded by its tier of evidence. The recurring result is that a smaller body is not a free good, and that its value depends on which accounts are kept and which are hidden.

Evidence Status

Established physiology. The budgets use the standard physiology of Part II. Demonstrated engineering. The components used in the model exist individually. Proposed design. The minimum habitable body is a design exercise, and each submission is a proposal to be assessed on its stated conditions.

Exercise 31.1. Carry out the capstone for a person whose specification includes a month in a hot climate with weekly access to service. State the binding constraint on the design and the stage of the transition path at which it would first appear.

Appendix A

Transport and Separation

This appendix collects the derivations used in Chapters 3, 6, and 18.

A.1 Series Resistance

For steady one-dimensional transport through layers $k = 1, \dots, N$ with flux J per unit area, the drop in potential across layer k is $\Delta c_k = J R_k$ with $R_k = \ell_k / (D_k \phi_k)$. Summing over layers gives $c_{\mathrm{in}} - c_{\mathrm{out}} = J \sum_k R_k$, which is Equation (3.2). Where the potential is discontinuous at an interface because of partitioning, the partition coefficient is absorbed into ϕ_k.

A.2 Minimum Work of Separation

Let a dilute ideal solution of osmotic pressure π_0 and volume V_0 be separated into pure solvent and a more concentrated retentate by withdrawing permeate volume V_p. When volume V_p has been withdrawn, the retentate has volume $V_0 - V_p$ and osmotic pressure $\pi = \pi_0 V_0 / (V_0 - V_p)$, since the amount of solute is fixed. The reversible work is

$$W = \int_0^{V_p} \pi \, \mathrm{d}V' = \pi_0 V_0 \ln \frac{V_0}{V_0 - V_p} = \pi_0 V_0 \ln \frac{1}{1 - r},$$

with $r = V_p / V_0$. Dividing by the permeate volume $r V_0$ gives Equation (6.3). For $r \to 0$ the expression tends to π_0.

A.3 Settling and Sedimentation Equilibrium

A sphere of radius r_p moving at velocity v through a fluid of viscosity μ experiences a drag of $6\pi\mu r_p v$ at low Reynolds number. Balancing this against the net body force $\tfrac{4}{3}\pi r_p^3 (\rho_p - \rho_f) a$ gives Equation (6.1). For a dissolved species at equilibrium in a centrifugal field, the chemical potential plus the potential energy in the field is uniform, and the condition $\mathrm{d}\mu / \mathrm{d}r = M(1 - \bar{v}\rho)\omega^2 r$ together with $\mu = \mu^0 + RT \ln c$ gives Equation (6.2).

A.4 Heat Conduction from a Sphere

Steady conduction from a sphere of radius R dissipating power P into an infinite medium of conductivity k gives a radial temperature profile $T(r) = T_\infty + P/(4\pi k r)$ for $r \geq R$, from which the surface temperature rise is $\Delta T = P/(4\pi k R)$, which is Equation (18.1). For uniform power density p_d the same result holds with $P = \tfrac{4}{3}\pi R^3 p_d$, giving Equation (6.4).

Perfusion adds a sink term and shortens the length over which the temperature decays, which lowers the rise.

Appendix B

Feedback and Stability

B.1 Loop with Delay

Consider the first-order process $\tau\dot{y} + y = Ku(t-\theta)$ under proportional control $u = -ky$. Taking Laplace transforms, the characteristic equation is $\tau s + 1 + kKe^{-s\theta} = 0$, which is Equation (7.1). At the stability boundary $s = j\omega$, so that $j\omega\tau + 1 + kK(\cos\omega\theta - j\sin\omega\theta) = 0$. The imaginary part gives $\omega\tau = kK\sin\omega\theta$ and the real part gives $1 + kK\cos\omega\theta = 0$. Eliminating kK gives $\tan\omega\theta = -\omega\tau$, that is $\omega\theta + \arctan(\omega\tau) = \pi$, and then $kK = \sqrt{1+\omega^2\tau^2}$. For $\theta \ll \tau$ the arctangent is close to $\pi/2$, so $\omega_c \approx \pi/(2\theta)$ and $(kK)_{\text{crit}} \approx \omega_c\tau = \pi\tau/(2\theta)$, which is Equation (7.2). To see that this is the first loss of stability, note that at $k = 0$ the only root is $s = -1/\tau$, in the left half-plane. The roots depend continuously on k, and for the retarded equation the number of roots with real part above any fixed level is finite, with none coming in from infinity. A root can therefore reach the right half-plane only by crossing the imaginary axis, and at $s = 0$ the equation gives $1 + kK = 0$, which has no solution for $kK > 0$. Crossings are thus at $s = j\omega$ with $\omega > 0$, and the smallest kK that satisfies the pair of equations above is the one at the smallest positive solution ω_c of $\omega\theta + \arctan(\omega\tau) = \pi$. Because the left side increases from 0 for $\omega > 0$, the solution is unique, and for $kK < \sqrt{1+\omega_c^2\tau^2}$ all roots stay in the left half-plane.

B.2 Saturating Actuators

Physiological actuators saturate: sweat rate has a maximum, vasodilation has a limit, and pumps have a maximum output. A loop that is stable for small disturbances can fail for large ones when the actuator saturates, because the effective gain falls to zero and the state drifts. The viable set $\mathcal{K}$ of Chapter 4 must therefore be checked against the saturation limits, and a design is acceptable only if the saturated capacity exceeds the largest disturbance in $\mathcal{D}$, which is the reserve condition of Equation (4.1).

B.3 Competing Controllers

When two controllers act on one variable, the net gain is the sum of their gains if both act in the same direction and can be near zero if they oppose one another. A device that regulates a quantity already regulated by the body should therefore be designed as a supplement to the native controller, with its action added to the native action in a known direction, or else it should override the native controller explicitly. The failure described in Worked Example 7.1 arises when the device acts on a path that is also used by the native controller and neither sees the other.

Appendix C

Reliability Models

C.1 The k-out-of-n Model

For n nodes each working independently with probability r, the number working X is binomial, and the organ is adequate when $X \geq k$. The probability is Equation (8.3). When n is large, the normal approximation with mean nr and variance $nr(1-r)$ gives $R \approx \Phi\big((nr - k + \tfrac{1}{2})/\sqrt{nr(1-r)}\big)$, with Φ the standard normal distribution function. The approximation shows that for $k/n < r$ the reliability approaches one as n grows, and for $k/n > r$ it approaches zero.

C.2 Heterogeneous Quality

If r is drawn from a Beta distribution with parameters a and b, the number of working nodes is beta-binomial with

$$\mathbb{P}(X = j) = \binom{n}{j} \frac{B(j + a,\, n - j + b)}{B(a, b)},$$

where B is the beta function. The mean is $na/(a+b)$ and the variance exceeds that of the binomial with the same mean by a factor that grows with n. The correlation between nodes is $1/(a+b+1)$, which for $a+b=10$ is about 0.09 and for $a+b=100$ about 0.01, a measure of the shared quality in Worked Example 8.1.

C.3 Common Cause

A shared layer that fails with probability q and disables all nodes gives Equation (8.4). When several shared layers are in series, their reliabilities multiply, and the reliability of the whole is bounded above by the product, which is the largest value that adding nodes could achieve.

C.4 Marginal Node

The increment $R_{n+1} - R_n = r\mathbb{P}(X_n = k-1)$ of Proposition 8.1 follows by conditioning on X_n. If $X_n \geq k$ the system is adequate with or without the new node. If $X_n \leq k-2$ it is inadequate in both cases. If $X_n = k-1$ the new node decides the outcome, and it works with probability r. For heterogeneous nodes the same argument holds with r replaced by the reliability of the added node.

C.5 Coverage by a Poisson Process

If node positions form a homogeneous Poisson process of intensity λ and each covers a ball of volume V_c, the number of nodes within reach of a fixed point is Poisson with mean λV_c. The probability that none is within reach is $e^{-\lambda V_c}$, which gives Equation (8.5). The expected uncovered fraction of a region equals this probability. For connectivity, the criterion is the percolation threshold of the continuum model, for which the critical mean degree of a three-dimensional random geometric graph is near 2.7 [31].

Appendix D

The Mathematics of Feasible Operating Ranges

D.1 Viability

Let $\dot{x} = f(x, u, d)$ be the dynamics of the person and devices, with control $u \in U$ and disturbance $d \in \mathcal{D}$. The viability kernel of a constraint set $\mathcal{K}$ is the set of initial states from which some admissible control keeps the state in $\mathcal{K}$ for all disturbances in the class and for all times in $[0, T]$ [2]. A specification is satisfiable at x_0 exactly when x_0 lies in the kernel.

D.2 Grace Period

Let u denote a support and $\mathcal{K}$ the viable set. For a state x_0 in the viability kernel under full support, the grace period for support u is

$$\tau_g(x_0; u) = \sup\{\tau \geq 0 : \ \exists \,\text{fallback control with } x(t) \in \mathcal{K} \quad \forall t \in [0, \tau],\ u \text{ withdrawn at } t = 0\}.$$

For the linear accumulation models of this book, τ_g reduces to a ratio of the remaining tolerable stock $D_{\lim} - D_0$ to a net rate, as in Equations (20.2) and (14.1), and it is zero when $D_0 \geq D_{\lim}$.

D.3 Intersections and Pareto Sets

If an architecture must satisfy several functions, its joint operating range is contained in the intersection of the ranges of the individual functions, and equals it under separability, as in Equation (30.1). A set of architectures can be compared by the pairs $(m(A), E(A))$ of carried mass and exposure. An architecture A dominates A' if $m(A) \leq m(A')$ and $E(A) \leq E(A')$ componentwise, with at least one strict inequality, and $R(A) \supseteq R(A')$, and the Pareto set is the set of architectures that no other dominates. The compaction problem of Equation (17.1) is solved when the Pareto set is found and the operating range of each member is stated.

D.4 Sensitivity of the Range to a Parameter

If the range is $R(A; \theta)$ for a parameter θ, such as a recovery fraction or a node reliability, the sensitivity of a bound is $\partial b / \partial \theta$ evaluated at the binding function when the minimum or maximum defining the bound is attained by a single function. Only that function has

a nonzero sensitivity of the bound, and at a tie the bound is not differentiable and only one-sided sensitivities exist. Away from ties, the improvement of any other function has no effect on the range at that bound. This property is the reason that Chapter 30 insists on identifying the binding constraint before attempting to improve a design.

Bibliography

[1] K. J. Åström and R. M. Murray. *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton University Press, 2008.

[2] J.-P. Aubin. *Viability Theory*. Birkhäuser, Boston, 1991.

[3] P. Bach-y-Rita, C. C. Collins, F. A. Saunders, B. White, and L. Scadden. Vision Substitution by Tactile Image Projection. *Nature*, 221:963–964, 1969.

[4] R. B. Bird, W. E. Stewart, and E. N. Lightfoot. *Transport Phenomena*. Wiley, revised 2nd edition, 2007.

[5] W. F. Boron and E. L. Boulpaep. *Medical Physiology*. Elsevier, 3rd edition, 2017.

[6] E. L. Cussler. *Diffusion: Mass Transfer in Fluid Systems*. Cambridge University Press, 3rd edition, 2009.

[7] P. Eckart. *Spaceflight Life Support and Biospherics*. Kluwer Academic Publishers and Microcosm Press, 1996.

[8] H. Flor, T. Elbert, S. Knöcht, C. Wienbruch, C. Pantev, N. Birbaumer, W. Larbig, and E. Taub. Phantom-Limb Pain as a Perceptual Correlate of Cortical Reorganization Following Arm Amputation. *Nature*, 375:482–484, 1995.

[9] Flyxion. Continuous Microsurgery: The Ethics of Never Putting the Knife Down. Manuscript, October 2026.

[10] Flyxion. Hepastitium and Relativistic Longevity. Unpublished monograph.

[11] Flyxion. Hyperpleonastic Legibility, section on Bimetallic Weaves as Mechanical Logic. Unpublished manuscript.

[12] Flyxion. Mechanical Inheritance. Unpublished research outline.

[13] Flyxion. Motile Womb Theory: State of the Research on Collagen and the Motile Womb. Research overview, 2026.

[14] Flyxion. Commit titled for Motile Womb Theory, repository standardgalactic.github.io, December 23, 2022. `https://github.com/standardgalactic/standardgalactic.github.io/commit/3e704f35c84eb7c3121dac923d69b63a5d604c67`. Earliest documented mention located.

[15] H. M. Frost. Bone "Mass" and the "Mechanostat": A Proposal. *The Anatomical Record*, 219(1):1–9, 1987.

[16] A. P. Gagge, J. A. J. Stolwijk, and Y. Nishi. An Effective Temperature Scale Based on a Simple Model of Human Physiological Regulatory Response. *ASHRAE Transactions*, 77(1):247–262, 1971.

[17] V. Gura, M. B. Rivara, S. Bieber, R. Munshi, N. C. Smith, L. Linke, J. Kundzins, M. Beizai, C. Ezon, L. Kessler, and A. Davenport. A Wearable Artificial Kidney for Patients with End-Stage Renal Disease. *JCI Insight*, 1(8):e86397, 2016.

[18] J. E. Hall and M. E. Hall. *Guyton and Hall Textbook of Medical Physiology*. Elsevier, 14th edition, 2021.

[19] M. A. Hofman. Evolution of the Human Brain: When Bigger Is Better. *Frontiers in Neuroanatomy*, 8:15, 2014.

[20] G. A. Ilizarov. The Tension-Stress Effect on the Genesis and Growth of Tissues, Part I: The Influence of Stability of Fixation and Soft-Tissue Preservation. *Clinical Orthopaedics and Related Research*, 238:249–281, 1989.

[21] M. Kleiber. Body Size and Metabolism. *Hilgardia*, 6(11):315–353, 1932.

[22] M. R. Mehra, Y. Naka, N. Uriel, D. J. Goldstein, J. A. Cleveland, P. C. Colombo, M. N. Walsh, C. A. Milano, C. B. Patel, S. Hutchins, J. Ransom, G. A. Ewald, S. J. Itoh, A. N. Salerno, F. D. Bates, J. Zucker, and C. J. Mullens. A Fully Magnetically Levitated Circulatory Pump for Advanced Heart Failure. *New England Journal of Medicine*, 376(5):440–450, 2017.

[23] B. Mota and S. Herculano-Houzel. Cortical Folding Scales Universally with Surface Area and Thickness, Not Number of Neurons. *Science*, 349(6243):74–77, 2015.

[24] N. C. Nowlan, J. Sharpe, K. A. Roddy, P. J. Prendergast, and P. Murphy. Mechanobiology of Embryonic Skeletal Development: Insights from Animal Models. *Birth Defects Research Part C*, 90(3):203–213, 2010.

[25] A. A. Pitsillides. Early Effects of Embryonic Movement: "A Shot out of the Dark". *Journal of Anatomy*, 208(4):417–431, 2006.

[26] M. Rausand and A. Høyland. *System Reliability Theory: Models, Statistical Methods, and Applications*. Wiley, 2nd edition, 2004.

[27] K. Schmidt-Nielsen. *Scaling: Why Is Animal Size So Important?* Cambridge University Press, 1984.

[28] R. Sender, S. Fuchs, and R. Milo. Revised Estimates for the Number of Human and Bacteria Cells in the Body. *PLoS Biology*, 14(8):e1002533, 2016.

[29] F. H. Silver, J. W. Freeman, and D. DeVore. Viscoelastic Properties of Human Skin and Processed Dermis. *Skin Research and Technology*, 7(1):18–23, 2001.

[30] S. Timoshenko. Analysis of Bi-Metal Thermostats. *Journal of the Optical Society of America*, 11(3):233–255, 1925.

[31] S. Torquato. *Random Heterogeneous Materials: Microstructure and Macroscopic Properties*. Springer, 2002.

[32] D. C. Van Essen. A Tension-Based Theory of Morphogenesis and Compact Wiring in the Central Nervous System. *Nature*, 385:313–318, 1997.

[33] S. Weinstein. Intensive and Extensive Aspects of Tactile Sensitivity as a Function of Body Part, Sex, and Laterality. In D. R. Kenshalo, editor, *The Skin Senses*, pages 195–222. Charles C. Thomas, Springfield, 1968.

[34] G. B. West, J. H. Brown, and B. J. Enquist. A General Model for the Origin of Allometric Scaling Laws in Biology. *Science*, 276:122–126, 1997.

[35] K. Zhang and T. J. Sejnowski. A Universal Scaling Law Between Gray Matter and White Matter of Cerebral Cortex. *Proceedings of the National Academy of Sciences*, 97(10):5621–5626, 2000.

[36] K. Zilles, N. Palomero-Gallagher, and K. Amunts. Development of Cortical Folding During Evolution and Ontogeny. *Trends in Neurosciences*, 36(5):275–284, 2013.

[37] J. Campisi and F. d'Adda di Fagagna. Cellular Senescence: When Bad Things Happen to Good Cells. *Nature Reviews Molecular Cell Biology*, 8:729–740, 2007.

[38] Flyxion. *The Architecture of Autonomy*. Manuscript. `https://github.com/standardgalactic/alphabet/blob/58037cac81cabccfdb968dc4469add455b527243/biology/architecture-of-autonomy.md`.

[39] Flyxion. *Closure of Closures*. Manuscript. `https://github.com/standardgalactic/alphabet/blob/58037cac81cabccfdb968dc4469add455b527243/biology/closure_of_closures.tex`.

[40] Flyxion. *Life Extension Suggestions*. Archived notes. `https://github.com/standardgalactic/kitbash/blob/5026f8f03eda99d0f3138ad95da01ba0b8d2a7ea/controversial-topics/background/Life%20extension%20suggestions.txt`.

[41] Flyxion. *Physiome: A Whole-Body Physiology Simulator in Rust*. Software repository. `https://github.com/standardgalactic/physiome`.

[42] Flyxion. *Recursive Continuation*. Manuscript. `https://github.com/standardgalactic/alphabet/blob/58037cac81cabccfdb968dc4469add455b527243/example/recursive-continuation.tex`.

[43] Flyxion. *The Geometry of Cellular Possibility*. See the section "Rejuvenation as Recovery of Admissible Futures." `https://standardgalactic.github.io/history/geometry-of-cellular-possibility.pdf`.

[44] Flyxion. *Restoration Before Explanation: Why Systems Can Be Repaired Before They Are Understood*. Manuscript, July 2026. `https://github.com/standardgalactic/calculus/blob/5ce306f7e687ca44eaf4db1d5666de8f9c183bc5/processing/restoration-before-explanation.tex`.

[45] Flyxion. *A Relativistic Theory of Longevity*. Manuscript. `https://github.com/standardgalactic/laboratory/blob/710bcc0a22892c07cd67a95743ae73ee49ad2857/source/relativistic-longevity.tex`.

[46] Flyxion. *Topological Therapeutics*. Manuscript. `https://github.com/standardgalactic/alphabet/blob/58037cac81cabccfdb968dc4469add455b527243/biology/topological-therapeutics.md`.

[47] Flyxion. *Xylomorphic Computation*. Archived transcript. `https://github.com/standardgalactic/abraxas/blob/4f4a3b4f6d63774d3192e6ae2a62f5092f7167ad/Xylomorphic_Computation.txt`.

[48] A. C. Guyton, T. G. Coleman, and H. J. Granger. Circulation: Overall Regulation. *Annual Review of Physiology*, 34:13–46, 1972.

[49] R. L. Hester, A. J. Brown, L. Husband, R. Iliescu, D. Pruett, R. Summers, and T. G. Coleman. HumMod: A Modeling Environment for the Simulation of Integrative Human Physiology. *Frontiers in Physiology*, 2:12, 2011.

[50] T. B. L. Kirkwood. Evolution of Ageing. *Nature*, 270:301–304, 1977.

[51] H. Kitano. Biological Robustness. *Nature Reviews Genetics*, 5:826–837, 2004.

[52] M. Levin. Bioelectric Signaling: Reprogrammable Circuits Underlying Embryogenesis, Regeneration, and Cancer. *Cell*, 184(8):1971–1989, 2021.

[53] J. P. Merrill, J. E. Murray, J. H. Harrison, and W. R. Guild. Successful Homotransplantation of the Human Kidney Between Identical Twins. *Journal of the American Medical Association*, 160(4):277–282, 1956.

[54] C. D. Murray. The Physiological Principle of Minimum Work: I. The Vascular System and the Cost of Blood Volume. *Proceedings of the National Academy of Sciences*, 12(3):207–214, 1926.

[55] M. Scheffer, J. Bascompte, W. A. Brock, V. Brovkin, S. R. Carpenter, V. Dakos, H. Held, E. H. van Nes, M. Rietkerk, and G. Sugihara. Early-Warning Signals for Critical Transitions. *Nature*, 461:53–59, 2009.

[56] N. Verzijl, J. DeGroot, S. R. Thorpe, R. A. Bank, J. N. Shaw, T. J. Lyons, J. W. J. Bijlsma, F. P. J. G. Lafeber, J. W. Baynes, and J. M. TeKoppele. Effect of Collagen Turnover on the Accumulation of Advanced Glycation End Products. *Journal of Biological Chemistry*, 275(50):39027–39031, 2000.

[57] J. Blow. If AI Solved Coding, Why Is My Software Still the Same. Video Interview, Transcript Consulted.

[58] D. Dixon. *Man After Man: An Anthropology of the Future*. Blandford, London, 1990.