Describing Multiple Levels of Abstraction in the Metabolism

Michael L. Mavrovouniotis

Department of Chemical Engineering
Northwestern University , Evanston, IL 60208-3120. U.S.A.
mlmavro@nwu.edu

Abstract

We discuss some central issues that arise in the computer representation of the metabolism and its subsystems. We provide a framework for the representation of metabolites and bioreactions at multiple levels of detail. The framework is based on defining an explicit linear mapping of metabolites and reactions from one level of detail to another. A simple reaction mechanism serves as an illustration and shows the emergence of the concept of a catalyst from metabolic abstraction levels.

Introduction

Many efforts have recently been undertaken to construct models, databases, and computer representations of the metabolism (Reddy *et al.*, 1993, Ochs and Conrow, 1991, Karp and Mavrovouniotis, 1994, Hunter, 1993, and references therein). The computational analysis of complex biochemical systems affects applications ranging from the development and improvement of industrial bioprocesses to the support of experimental studies of cell physiology (Karp and Mavrovouniotis, 1994).

Emerging new representations of biochemical pathways and biochemical knowledge-bases allow the storage and retrieval of commonly available types of knowledge and data (for biochemical compounds, reactions, pathways, and enzymes); the representations are intended to be flexible enough to accommodate many types of automated reasoning. Through these efforts, a transition is under way from narrow and specialized approaches towards comprehensive general-purpose computer models of the metabolism. Flexible multi-purpose representations are obviously important in avoiding duplication of effort (since there is often significant overlap in the information needed by many different applications or problems). The flexibility is important even in a single target problem, because one is often forced to start solving parts of the problem before finalizing the overall problem formulation and solution strategy, i.e., one develops programs initially in small chunks, which are extended and connected later. Within this context, the underlying computational representations must be created, from the outset, general enough to support the needs of all the data, algorithms, and display capabilities; these might not be clear at the outset, because one usually begins from narrow formulations and simple examples.

Many obstacles and questions stand in the way of general and flexible computational representations of the metabolism. The issue that will be discussed in this paper is the representation of biochemical systems at multiple levels of abstraction, which are always present (often informally or implicitly) in any description of a biochemical system. As is often the case, computer representation efforts (and in particular the assignment of semantics to computational models or knowledge-bases) bring forth terminological ambiguities that were there to start with. Thus, the impact of formalization extends to the framework of the domain itself, beyond the development of algorithms and software.

We focus on ambiguities that are a direct consequence of multiple levels of detail in biochemical representations. The enzymes that catalyze bioreactions and the compounds which participate in the reactions may have variants in their structure, because of the effects of water, the formation of complexes, and regulatory modifications. The variants may be important in one level of detail but should be ignored in another. How do we build the connection between a lumped and a detailed view? Bioreactions are often organized into pathways or decomposed into steps; what are the implicit criteria used in this decomposition? How can this decomposition be made explicit in the computer representation of this hierarchical system? This paper aims to discuss and clarify these key issues that arise in the formalized or computational representation of the metabolism at multiple levels of detail.

The remainder of the introduction expands on the biological motivation of the work and introduces the basic terminology and assumptions we will use. Subsequent sections present the concept of a metabolic view (which defines a specific level of detail) and the abstraction of one metabolic view to a less detailed metabolic view.

Ambiguity and Abstraction in the Metabolism

The biochemical transformations that comprise the metabolism have a hierarchical structure. What we normally think of as one bioreaction is actually composed from several elementary steps, which are often called, collectively, the mechanism of the reaction. The

This work was supported in part by the National Library of Medicine (grant R29 LM05278).

mechanism includes association steps that form complexes of the enzyme and the substrates, steps that actually alter the chemical structure of the substrates, and dissociation steps that release products from the enzyme complex.

What is an enzymatic reaction, and where does it begin and end? If we have two enzymes which are attached to each other, shuttling an intermediate directly between them, should that count as one reaction or as two? If the product of an enzyme-mediated reaction remains bound to the enzyme, should the reaction be considered complete? The difficulty in answering such questions belies the informal way in which we have abstracted a multitude of intra- and intermolecular phenomena into higher level concepts such as "enzymatic reaction" or "pathway".

How much significance do we attribute to assemblages such as pathways, and their a priori designation? The computerized metabolic map of Ochs and Conrow (1990) attempts to hardwire the organization of bioreactions into pathways. Mavrovouniotis *et al.* (1990) attribute little significance to *a priori* designation of pathways; they construct pathways as bioreaction combinations which achieve a target transformation, regardless of whether they are designated as such in a biochemistry textbook. In either case, pathways are abstractions, and their relationship to the underlying detailed biochemical network should be made explicit.

The definition of bioreactions and pathways from the soup of interacting molecular structures represents a hierarchical abstraction created by *choice*, rather than compelling physical separation. The isolation of one subset of an enzyme's interactions as a biochemical reaction, and another subset as a regulatory effect, are not dictated by physical laws. There are good reasons for conforming to this hierarchy, but the fact that it is not compelling means that it is not completely predefined. For non-computational representations some ambiguity (and even the occasional mild conceptual inconsistency) can be useful, but in computational knowledge representation the hierarchy must be defined explicitly and consistently.

The thrust of this paper is the modeling of multi-level abstraction in the metabolism. Abstraction is invoked whenever we describe an entire set of species (such as acid or base forms, or different enzyme-metabolite complexes) as a single entity. Abstraction is also involved when many physicochemical events (such as formation, reaction, and dissociation of enzyme-substrate complexes) are lumped into a single process ("enzymatic reaction"). Even the meaning of a term like "catalyst" is a result of abstraction, as we will discuss in the last section of this paper. We propose a framework that makes these abstractions explicit, allowing formal systems or computer implementations to model, carry out, and reason about such abstractions.

Preliminaries of Abstraction Scheme

We begin with some essential background definitions, assumptions, and conventions which will facilitate the presentation of the multi-layer abstraction scheme. For simplicity, we will assume that we are modeling a closed system (a system in which there is no influx or outflow of material); this means that all changes in amounts of metabolites are due to reactions. The extension to open systems is conceptually easy. This simplification is intended to avoid detractions from the core issues.

We will be referring often to amounts of metabolites. Their measurement units will always be moles; the number of moles is proportional to the number of molecules (with Avogadro's number 6.02×10^{23} as the proportionality constant). For a constant-volume system we can substitute concentrations (in mol/l=kmol/m^3=M) for amounts. We are primarily interested in *changes* in amounts, rather than absolute values.

The amounts of chemical elements present in a system will be treated like the amounts of metabolites, above. What is important about chemical elements is that they are conserved by reactions (while metabolite amounts obviously are not). Thus, we can use as elements any kind of conserved substructures or moieties; if aromatic rings are unaffected by the reactions in our system, then we can treat an aromatic ring like an element. The terms element and moiety will be used interchangeably for conserved substructures.

To describe the quantitative progress of a bioreaction, we use the reaction *extent*. It is expressed in moles (or concentration units). We can think of a change in reaction extent as a measurement of how many times the reaction has occurred at the molecular level, translated to macroscopic units (from molecules to moles) through division by Avogadro's number. For a closed system, the time-derivative of the reaction extent is the reaction rate.

We will assume that biotransformations are reversible, so that the reaction extents can be positive or negative. As we have already begun to do, we will use the terms biotransformation, bioreaction, or reaction interchangeably.

The following basic linear algebra concepts will be used:

- The contents of matrices will be enclosed in brackets. The entry at the intersection of row i and column j of a matrix W will be indicated as $W(i,j)$. The transpose of a matrix W will be indicated as W^T. A row vector will have the form $[v_1\ v_2\ v_3\ \ldots]$; a column vector can be written as the transpose of a row vector $[v_1\ v_2\ v_3\ \ldots]^T$.
- Identity matrices will be indicated by I and zero matrices (or vectors) by 0; they will be assumed to have appropriate dimensions for the equation in which they occur.
- The columnspace of a matrix W is the space of all vectors that are linear combinations of the columns of W, i.e., all vectors that are equal to Wv for some v. The rowspace of W is the columnspace of W^T. The nullspace of W is the set of vectors v such that Wv=0. The left nullspace is the nullspace of W^T.
- The pseudoinverse (Strang, 1986) or Moore-Penrose inverse of W will be denoted as W^+.

We have opted for a matrix view because it makes the notation compact. In actual implementation, sparse matrices often take the form of symbolic representations. In an object-oriented scheme, the matrices would in all

likelihood be implicit in attributes of objects, and matrix operations would translate to iterations over attributes and objects. The matrix notation used here does not require (and does not propose) the use of arrays to represent relations in the metabolism.

Metabolic Views

In the introduction, we discussed the question of what constitutes a distinct biochemical compound and what simply an alternative form of the same compound. In the same spirit, we discussed the distinctions between bioreactions and their steps and between pathways and bioreactions. We propose here a multi-layer scheme for making these distinctions consistently at several levels of abstraction. We will illustrate our proposal with a simple example as we go along.

A metabolic view describes a chosen level of detail; it consists of three types of entities (metabolites, elements, and bioreactions), whose amounts provide the quantitative description of changes in the system's state. Linear mappings define quantitative relations among these entities' amounts. Constraints among the linear mappings ensure the internal consistency of the metabolic view.

Entities

There are three types of basic building blocks for a metabolic view.

- A set of n_b metabolites, $B_1, B_2, \ldots$, whose ordering (and indices) will be used to condense, into matrices, other information about them. Metabolite amounts, or changes in their amounts, will be represented by a vector b ($n_b \times 1$). The signs of the components of b are not restricted.
- A set of n_y conserved elements or moieties, $Y_1, Y_2, \ldots$, that will be used to describe the molecular formula or structure of the metabolites; the ordering and indices of the moieties will be used to condense information into matrices. Element amounts, or changes in their amounts, will be represented by a vector y ($n_y \times 1$).
- A set of n_p biotransformations or bioreactions, $P_1, P_2, \ldots$, that take place among the metabolites; the indices will likewise be used for encoding information into matrices. Changes in reaction extents will be represented by a vector p ($n_p \times 1$). The vector p may have negative entries, since reactions are reversible (and we are not restricting p to represent changes in a specific time sequence).

Example 1

We consider the mechanism of a single-substrate, single-product enzymatic reaction. The metabolites B_1 and B_2 are the substrate and product, B_3 is the free enzyme, and B_4 and B_5 are enzyme-substrate and enzyme-product complexes. If a system's composition changes by 1 mol of B_5, -1 mol of B_1 and -1 mol of B_3 (i.e., there is an increase in B_5, but a decrease in B_1 and B_3) the corresponding vector b would be written as:

$b=[-1\ 0\ -1\ 0\ 1]^T$

There are three reactions:

P_1: $B_1 + B_3 \rightarrow B_4$

P_2: $B_4 \rightarrow B_5$

P_3: $B_5 \rightarrow B_2 + B_3$

If reaction P_1 advances by 1 mol and reaction P_3 by -1 mol (i.e., 1 mol in the reverse direction), the corresponding vector p would be written as $p=[1\ 0\ -1]^T$.

To describe the elemental composition, we could refer to the molecular formulas of the species, and use the elements that occur in it. For example, if B_1 is dihydroxyacetone phosphate, B_2 is glyceraldehyde phosphate, and B_3 is the enzyme *glyceraldehyde phosphate dehydrogenase* (from a specific source), we could use as elements C, H, O, N, and S (the last two needed only for B_3, B_4, and B_5). However, any set of conserved moieties is acceptable, and we can avoid unnecessary detail by using a small set of large moieties. In this case, we can manage with just two conserved moieties, i.e., the moiety Y_1 for the substrate and product (B_1 and B_2), and the free-enzyme moiety Y_2; only the complexes B_4 and B_5 contain both moieties.

We have, in summary, $n_b=5$, $n_y=2$, and $n_p=3$. We will intersperse this example with the exposition of the abstraction framework.

Linear Mappings

The vectors b, y, and p that describe a change in a metabolic system are related to each other. A change in p (bioreaction extents) should clearly map to a change in b (metabolites), and a change in b should map to a change in y. We will restrict these mappings to be linear, because the two properties defining a linear transformation are essential in the use of these quantities. Specifically, if p maps to b, then Kp should map to Kb, enforcing an essential notion of proportionality. If furthermore p' maps to b' then p+p' should map to b+b'; this allows two successive changes in a system to be either modeled individually or consistently lumped into one. The same arguments can be made for the other mappings we use here. Thus, they are linear mappings that can be represented by matrices.

Three linear mappings define quantitative the necessary relations among building blocks.

- A linear transformation from metabolites to elements translates amounts (or changes in amounts) of metabolites into amounts (or changes in amounts) of elements. This transformation is represented by a matrix E_b ($n_y \times n_b$) describing the composition of each metabolite (column) in terms of elements. The $E_b(i,j)$ entry in the matrix is the number of occurrences of element Y_i in metabolite B_j. A change of b in metabolite amounts translates into a change of y in element amounts:

$$y = E_b\, b$$

In the ordinary case, moieties Y_i represent actual chemical elements and metabolites B_j represent actual chemical compounds; the entries of E_b would be non-negative in this case (with the possible exception of a Y_i representing electrical charge). However, in our study of multi-layer abstraction, moieties may take other forms, and it may be necessary to model a metabolite as the *difference* of moieties. Thus, E_b is permitted to have negative entries.

• A linear transformation from reactions to metabolites translates changes in biotransformation extents into changes in amounts of metabolites. This transformation is represented by a matrix R_{bp} ($n_b \times n_p$) describing the stoichiometry of each biotransformation (column) in terms of metabolites. The $R_{bp}(i,j)$ entry in the matrix is the stoichiometric coefficient of metabolite B_i in bioreaction P_j; it is positive if B_i is a product and negative if B_i is a reactant of P_j. A change of p in bioreaction extents translates into a change of b in metabolite amounts:

$$b = R_{bp}\, p$$

• A linear transformation from metabolites to reactions translates changes in amounts of metabolites into changes in biotransformation extents. This linear transformation, which is represented by a matrix S_{bp} ($n_p \times n_b$), is intended to provide the system-identification aspects of biotransformations. Clearly the occurrence of a biotransformation is the cause, the internal mechanism that gives rise to the changes in metabolite amounts. The transformation S_{bp} maps observed effects (measurements of changes in metabolite-amounts) to changes in internal causes (reaction extents). A change of b in metabolite amounts translates into a change of p in bioreaction extents:

$$p = S_{bp}\, b$$

The general need, throughout this framework, to define forward and reverse transformations explicitly is a direct consequence of the fairly unrestricted shape and content of the matrices. The matrix R_{bp} is not restricted to be square; it may be either fat ($n_b < n_p$) or tall ($n_b > n_p$), and it may have dependent rows and/or dependent columns. As a consequence, no simple definition of an inverse S_{bp} from R_{bp} will suffice for all cases.

Example 1, continued

The composition of the chemical species in terms of moieties reflects our choice of moieties. The metabolites B_1 and B_2 contain only Y_1, B_3 contains only Y_2, while B_4 and B_5 contain both Y_1 and Y_2:

$$E_b = \begin{bmatrix} 1 & 1 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 \end{bmatrix}$$

If a system's composition changes by $b=[-1\ 0\ -1\ 0\ 1]^T$, then the moieties would change by $y=E_b b$, or:

$$y = \begin{bmatrix} 1 & 1 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} -1 \\ 0 \\ -1 \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

If the change in the composition is instead $b=[1\ 0\ 0\ 0\ 0]^T$, then $y=[1\ 0]^T$.

At this level of abstraction, the matrix of stoichiometries of the three reactions is:

$$R_{bp} = \begin{bmatrix} -1 & 0 & 0 \\ 0 & 0 & 1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \\ 0 & 1 & -1 \end{bmatrix}$$

A change of $p=[1\ 0\ -1]^T$ in bioreaction extents translates into a change of b in metabolite amounts $b=R_{bp}p=[-1\ -1\ -2\ 1\ 1]^T$. For mapping metabolite changes to reactions, we will use

$$S_{bp} = \frac{1}{8} \begin{bmatrix} -5 & 3 & -2 & 1 & 1 \\ -4 & 4 & 0 & -4 & 4 \\ -3 & 5 & 2 & -1 & -1 \end{bmatrix}$$

A change of b in metabolite amounts $b=[-1\ -1\ -2\ 1\ 1]^T$ translates back to bioreaction extents $p=S_{bp}b=[1\ 0\ -1]^T$. A change of $b=[1\ 0\ 0\ 0\ 0]^T$ can clearly not be derived from any choice of bioreaction extents, i.e., there is no p such that $b=R_{bp}p$. But the mapping S_{bp} nevertheless produces a corresponding $p=S_{bp}b=[-5/8\ -1/2\ -3/8]^T$.

Constraints

The intended physical meaning of the mappings introduces constraints for the matrices.

• A change p in bioreaction extents translates into $b=R_{bp}p$ and $y=E_b b=E_b R_{bp}p$. Elements Y_i are intended to be conserved by the biotransformations. Thus, for any p, we must have $E_b R_{bp}p=0$. Therefore:

$$E_b\, R_{bp} = 0$$

• The relation between S_{bp} and R_{bp} is less clear. The matrices are conceptual, but not mathematical, inverses. There are two extreme cases. The first case occurs when the columnspace of the matrix R_{bp} spans the nullspace of E_b; then, any change in metabolites that conserves moieties can be brought about through reactions present in the system. For this case an invertibility constraint between S_{bp} and R_{bp} takes the following form: For all b in the nullspace of E_b, we must have $R_{bp}S_{bp}b=b$ (i.e., an element-conserving b should map reversibly to reactions). The second case is when R_{bp} is full column rank, i.e., the reaction stoichiometries are linearly independent; then, the reaction extents are identifiable, and the suitable constraint is $S_{bp}R_{bp}=I$. These two cases are neither mutually exclusive nor exhaustive.

Example 1, continued

We can easily verify that the constraint $E_b R_{bp}=0$ holds, i.e., the reactions are consistent with the conservation of moieties.

This example falls into one of the two special cases, with respect to the relation between R_{bp} and S_{bp}: The matrix R_{bp} is full column rank. We can verify that $S_{bp}R_{bp}=I$. Therefore, given any reaction extents p, if we transform it to the corresponding metabolite amounts $b=R_{bp}p$, then the transformation back to reaction extents $S_{bp}b$ will yield the original p. On the other hand:

$$R_{bp}S_{bp} = \frac{1}{8} \begin{bmatrix} 5 & -3 & 2 & -1 & -1 \\ -3 & 5 & 2 & -1 & -1 \\ 2 & 2 & 4 & -2 & -2 \\ -1 & -1 & -2 & 5 & -3 \\ -1 & -1 & -2 & -3 & 5 \end{bmatrix}$$

which exhibits large diagonal elements but is not equal to I. However, the columnspace of R_{bp} does span the nullspace of E_b. Therefore, starting with metabolite amounts $b=[-1\ 1\ 0\ 0\ 0]^T$ for which $E_b b=0$, if we map to reactions $p=S_{bp}b=[1\ 1\ 1]^T$ and then back to metabolites $R_{bp}p$ we recover the original b.

Strict Abstraction of Metabolic Views

Let metabolic view B be characterized by n_b metabolites

$B_1, B_2, \ldots$, with amount changes described by b ($n_b \times 1$); n_y conserved elements, $Y_1, Y_2, \ldots$, with changes described by y ($n_y \times 1$); and n_p bioreactions, $P_1, P_2, \ldots$, with changes in reaction extents described by p ($n_p \times 1$). The mappings are denoted by the matrix E_b ($n_y \times n_b$) for the composition of metabolites in terms of elements; the matrix R_{bp} ($n_b \times n_p$) for the stoichiometries of bioreactions; and the matrix S_{bp} ($n_p \times n_b$) for mapping changes in metabolite-amounts to changes in reaction extents. We assume that the constraints on the matrices are satisfied.

For the same underlying metabolic system, we may have another metabolic view, A. We will accept that the n_y conserved elements, $Y_1, Y_2, \ldots$, are the same, as a first precondition for a relation between A and B. The remaining building blocks may differ. Let A involve n_a metabolites $A_1, A_2, \ldots$, with amount changes described by a ($n_a \times 1$); and n_q bioreactions, $Q_1, Q_2, \ldots$, with changes in reaction extents described by q ($n_q \times 1$). The mapping matrices are denoted by E_a ($n_y \times n_a$), R_{aq} ($n_a \times n_q$), and S_{aq} ($n_q \times n_a$), and they are assumed to satisfy the necessary constraints.

We will use the symbol $\in$ to identify choices made from view A or view B. Thus, $a \in A$ denotes a vector (with dimensions $n_a \times 1$) of metabolite amounts from view A, while $q \in A$ denotes a vector of bioreaction extents from view A.

A *strict abstraction* of B to A is intended to model the notion that B is a detailed view of the metabolic system and A is a coarse view of the same system. We use the term strict to show that in all areas of the system the view B is at least as detailed as the view A (i.e., there is no part of the system in which A is more detailed than B).

The abstraction must include a way to interpret amounts (more importantly, changes in amounts) of metabolites from one view to another. The mapping of amounts across views must be linear, because we want proportionality and superimposable changes, as a system goes through a sequence of states: If a maps to b, then Ka should map to Kb, enforcing an essential notion of proportionality; if furthermore a' maps to b' then a+a' should map to b+b'. A strict abstraction of B to A is described by linear transformations, in matrix form, and a set of constraints.

Figure 1 shows the mappings of views A and B, along with the mappings that define the abstraction and will be presented below.

Example 1, continued

We postulate view A to eliminate the enzyme-substrate and enzyme-product complexes (in one of the many possible ways). We will have only three metabolites: A_1, intended to replace the substrate B_1; A_2, intended to replace the product B_2; and A_3, intended to replace the enzyme B_3. The elemental matrix E_a takes the form:

$$E_a = \begin{bmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

i.e., A_1 and A_2 contain only moiety Y_1, while A_3 contains only Y_2. The one reaction will take the form:

Q_1: $A_1 \rightarrow A_2$, with $R_{aq} = [-1\ 1\ 0]^T$, $S_{aq} = [-1/2\ 1/2\ 0]$

If we prespecify this reaction as the desired outcome, we must design the transformations (discussed below) so that the abstraction will accomplish this. Alternatively, we can select an intuitively appropriate set of transformations, and derive the reactions in A from the constraints presented later. We summarize the dimensions in view A as $n_a=3$, $n_q=1$, and $n_y=2$.

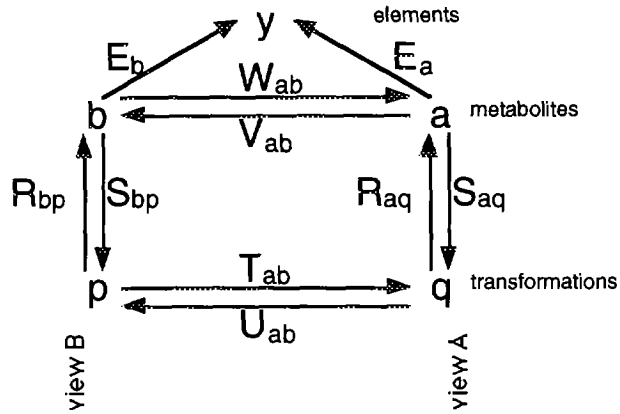

Figure 1. The transformations of the vectors describing two metabolic views.

Transformations of Metabolites

• A linear transformation of changes in metabolites from view B (vector b) to view A (vector a) is represented by matrix W_{ab} ($n_a \times n_b$), permitted to have negative entries. A change $b \in B$ translates into a change $a \in A$:

$a = W_{ab}\, b$

• A corresponding transformation of changes from view A to view B is represented by a matrix V_{ab} ($n_b \times n_a$). A change $a \in A$ translates into a change $b \in B$:

$b = V_{ab}\, a$

Example 1, continued

We adopt the following transformations for the abstraction of B to A:

$$W_{ab} = \begin{bmatrix} 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 1 & 1 \end{bmatrix}, \quad V_{ab} = \frac{1}{8}\begin{bmatrix} 5 & 1 & -2 \\ 1 & 5 & -2 \\ -2 & -2 & 4 \\ 3 & -1 & 2 \\ -1 & 3 & 2 \end{bmatrix}$$

The mapping of B_i amounts to A_i amounts is represented in matrix W_{ab}. Viewed by rows, W_{ab} defines metabolite A_1 as occurring in both B_1 and B_4; A_2 occurring in both B_2 and B_5; and enzyme A_3 occurring in B_3, B_4, and B_5. All of these occurrences are in 1:1 molar proportions. Viewing W_{ab} by columns shows that B_1 is treated as merely A_1; B_2 is treated as A_2; the species B_4 is treated by W_{ab} as an enzyme-substrate (A_3+A_1) complex; and the species B_5 is treated as an enzyme-product (A_3+A_2) complex. The matrix W_{ab} effectively contains a common view of the species.

But what of matrix V_{ab}? It translates A_i amounts to B_i amounts in a way that may seem peculiar, but there is no neat way for this transformation. We are transforming a low-dimensional

(or low information-content) vector into a higher dimensional one. If there is a change by 1 mol for A_3 (i.e., $a=[0\ 0\ 1]^T$), how should we interpret or allocate this change in terms of B_3, B_4, and B_5? With the chosen V_{ab}, we obtain $b=[-1/4\ -1/4\ 1/2\ 1/4\ 1/4]^T$; the change is allocated 1/2 to B_3, and 1/4 each to B_4, and B_5. There is even a change in B_1 and B_2, which has the beneficial effect of balancing moiety Y_2, since:

$$E_a[0\ 0\ 1]^T = [0\ 1]^T = E_b[-1/4\ \ -1/4\ \ 1/2\ \ 1/4\ \ 1/4]^T.$$

Constraints on Metabolites and Elements

The constraints that must be satisfied by W_{ab} and V_{ab} (in conjunction with E_a and E_b) derive from our notion of strict abstraction and conservation of elements. In this exposition, we will not attempt to separate the constraints into a core definition and a set of derived properties. Rather, we will focus on clarifying the abstraction process and include all the constraints that offer insight to this process.

• We first note that $n_b \geq n_a$ if B is at least as detailed as A. In the case of equality, we have square matrices W_{ab} and V_{ab}, which must be non-singular and true inverses ($W_{ab}V_{ab}=I=V_{ab}W_{ab}$). This is clearly a trivial abstraction, since A and B, although not identical views, represent the same level of detail. We are primarily interested in the case $n_b > n_a$, which makes W_{ab} fat and V_{ab} tall.

• Since B is the more detailed level, there are likely to be many different states in B that correspond to the same state in A, i.e., for any given vector $a \in A$, there are in general several vectors $b \in B$ with $a=W_{ab}b$. The only constraint we can rationally introduce, however, is that there should be at least one such vector $b \in B$; this ensures that all values in the more abstract view A have precise counterparts in the more detailed view B. This constraint can be stated in many equivalent forms, including that the matrix W_{ab} has a columnspace of dimension n_a (i.e., its columns span the space), that the matrix W_{ab} has an empty left nullspace, or that the rows of W_{ab} are linearly independent. We will opt for the equivalent statement:

$$\text{rank}(W_{ab})=n_a$$

• We expect that for any vector $b \in B$ there is at most one vector $a \in A$ with $b=V_{ab}a$. A suitable vector $a \in A$ will not exist for all $b \in B$, but different values of $a \in A$ should never produce the same value for $b \in B$. This constraint, like the previous one, can be stated in many equivalent forms: The matrix V_{ab} has a columnspace of dimension n_a (i.e., it has linearly independent columns); the matrix V_{ab} has an empty nullspace; or the rows of V_{ab} span the space. We opt for the statement:

$$\text{rank}(V_{ab})=n_a$$

• Suppose we select $a \in A$, we compute $b=V_{ab}a$, and then compute $W_{ab}b=W_{ab}V_{ab}a$. This means that we identified a detailed $b \in B$ that we match with the initial choice of $a \in A$, and then we went back to the abstract level. We should get back the originally chosen $a \in A$. In effect, among the many solutions (for a given $a \in A$) to $a=W_{ab}b$, the intent is for the matrix V_{ab} to select one ($V_{ab}a$), but not to produce a solution that relates to an altogether different $a \in A$. Consequently:

$$W_{ab}\,V_{ab} = I$$

• The above relation, a result of the view B being a refinement of view A, states that going from $a \in A$ to $b \in B$ and then back to $a \in A$ is an identity mapping. We cannot make this argument for the mapping going from $b \in B$ to $a \in A$ and then back to $b \in B$; the product $V_{ab}W_{ab}$ is generally *not* equal to the identity matrix. However, if $b \in B$ is in the column space of V_{ab}, then $V_{ab}W_{ab}b=b$. This can be easily seen from the fact that if b is in the columnspace of V_{ab} then $\exists\, a$ such that $b=V_{ab}a$, hence $V_{ab}W_{ab}b=V_{ab}W_{ab}(V_{ab}a)=V_{ab}(W_{ab}V_{ab})a=V_{ab}a=b$. If b lies in the left nullspace of V_{ab}, then $V_{ab}W_{ab}b$ will also lie in the left nullspace of V_{ab}; if b lies in the nullspace of W_{ab} then $V_{ab}W_{ab}b=0$. If we chose to go one step further and require the last two subspaces to coincide, then W_{ab} and V_{ab} would be Moore-Penrose inverses (or pseudoinverses) of each other; we will refrain from this restriction.

• What about the conservation of elements? In $b=V_{ab}a$ we are refining to a more detailed level, and we would certainly expect $E_bb=E_aa$. Therefore $E_bV_{ab}a=E_aa$ for any $a \in A$, or:

$$E_b\,V_{ab} = E_a$$

• The mapping $a=W_{ab}b$ cannot be guaranteed to conserve elements; generally, $E_aW_{ab} \neq E_b$. But if $b \in B$ lies in the columnspace of V_{ab}, then $E_bb=E_bV_{ab}a=E_aa=E_aW_{ab}b \Rightarrow E_aW_{ab}b=E_bb$.

Example 1, continued

The chosen matrix W_{ab} has the appropriate rank 3. Thus, any metabolite combination $a \in A$ can be produced by some $b \in B$, i.e., $W_{ab}b=a$ can always be solved for b. The chosen V_{ab} also has rank 3; the mapping of metabolite amounts from A to B will never produce the same b from different a's.

We already illustrated the constraint $E_bV_{ab}=E_a$, by showing that for $a=[0\ 0\ 1]^T$, $b=V_{ab}a=[-1/4\ \ -1/4\ \ 1/2\ \ 1/4\ \ 1/4]^T$, one obtains $E_a[0\ 0\ 1]^T=[0\ 1]^T=E_b[-1/4\ \ -1/4\ \ 1/2\ \ 1/4\ \ 1/4]^T$, i.e., $E_bV_{ab}a=E_aa$.

Matrix multiplication shows that $E_bV_{ab}=E_a$ and that $W_{ab}V_{ab}=I$. We have $V_{ab}W_{ab} \neq I$, although there is a clear pattern of large diagonal elements (because we in fact selected V_{ab} and W_{ab} as pseudoinverses):

$$V_{ab}W_{ab}=\frac{1}{8}\begin{bmatrix} 5 & 1 & -2 & 3 & -1 \\ 1 & 5 & -2 & -1 & 3 \\ -2 & -2 & 4 & 2 & 2 \\ 3 & -1 & 2 & 5 & 1 \\ -1 & 3 & 2 & 1 & 5 \end{bmatrix}$$

The mapping $a=W_{ab}b$ is usually not guaranteed to conserve elements. In this case, however, it happens to be conserving, i.e., $E_b=E_aW_{ab}$.

As we verify the validity of the stated constraints, and point out special properties of this example, one should bear in mind that the validity of the properties is not coincidental; it is a result of the careful choice of the view A and the mapping matrices.

Transformation of Reactions

Forward and reverse transformations are used for reaction extents, in a manner similar to metabolites.

• The matrix T_{ab} ($n_q \times n_p$) maps reaction extents from view B to view A.

$q = T_{ab}\, p$

• The matrix U_{ab} ($n_p \times n_q$) maps reaction extents from view A to view B.

$p = U_{ab}\, q$

Example 1, continued

We will set the matrices for the example as follows:

$T_{ab} = [0\ 1\ 0]$

$U_{ab} = [0\ 1\ 0]^T$

The extent of reaction $Q_1 \in A$ is taken simply as equal to the extent of reaction P_2. Reactions P_1 and P_3 do not feature at all in the mapping, for a simple reason: They are actually null reactions in view A! If we have $p=[1\ 0\ 0]^T$, i.e., 1 mol of reaction P_1, we obtain $b=[-1\ 0\ -1\ 1\ 0]^T$, which becomes $a=W_{ab}b=[0\ 0\ 0]^T$. This is a direct consequence of our choice to consider B_4 as a combination of A_1+A_3.

There is actually very little freedom left, once the view A and the metabolite matrices have been specified; through careful choice of T_{ab} and U_{ab}, we have essentially anticipated the constraints discussed below.

Constraints on Reactions and Metabolites

• In accordance with our view of abstraction, we expect $n_p \geq n_q$. Equality $n_p=n_q$ would mean that there are precisely as many reactions in the abstract view A as in the detailed view B, making the abstraction trivial. In the non-trivial case $n_p>n_q$, the matrix T_{ab} is fat and the matrix U_{ab} is tall.

• In a manner similar to the derivation of the relation between W_{ab} and V_{ab}, we require that

$T_{ab}U_{ab}=I$

• We must accept that in general $U_{ab}T_{ab} \neq I$, but for $p \in B$ lying in selected subspaces we will have $U_{ab}T_{ab}p=p$. We will not provide an analysis similar to that given for the case $V_{ab}W_{ab}b=b$. Note that the existence and the dimension of a subspace in which $U_{ab}T_{ab}p=p$ can be derived from the fact that if $T_{ab}U_{ab}=I$ the matrix $U_{ab}T_{ab}$ must have 1 as an eigenvalue.

• A change in reaction extents $p \in B$ leads to changes in amounts by $R_{bp}p$ which translate, in view A, to $W_{ab}R_{bp}p$. The same change $p \in B$ translates to a change $q \in A$, $q=T_{ab}p$, which means a change in metabolite amounts by $R_{aq}T_{ab}p$. We require that the two results be identical, $W_{ab}R_{bp}p=R_{aq}T_{ab}p$ for any p. Thus:

$W_{ab}R_{bp}=R_{aq}T_{ab}$

$\Rightarrow W_{ab}R_{bp}U_{ab}=R_{aq}T_{ab}U_{ab} \Rightarrow W_{ab}R_{bp}U_{ab}=R_{aq}$

• Following the above discussion, column i of $W_{ab}R_{bp}$ represents the reaction stoichiometry of the corresponding $P_i \in B$ translated into view A. In a strict abstraction, the reaction space of A derives entirely from the reaction space of B. The columns of the matrix R_{aq} should thus be a permutation of the columns of $W_{ab}R_{bp}$, with duplicates potentially eliminated and columns equal to zero removed. In effect, each non-zero column of $W_{ab}R_{bp}$ should be proportional to precisely one column of R_{aq}.

Thus, if we require that the reactions in A be no more and no less than the above translation of the reactions from B, we can restrict the matrix T_{ab} (from the first form of the previous equation) to have at most one non-zero element per column. Similarly, from the second form of the above equation, the matrix U_{ab} should have at most one non-zero element per row. Furthermore, if $T_{ab}(i,j)$ is the non-zero entry for column j of T_{ab}, then $U_{ab}(j,i)$ is the only entry permitted to be non-zero in row j of U_{ab}. In fact, if we did not mind having in A null reactions (zero stoichiometry) and duplicate (except for a proportionality factor) reactions, we could require $W_{ab}R_{bp}=R_{aq}$, setting $n_p=n_q$ and $U_{ab}=T_{ab}=I$.

Example 1, continued

We can easily verify that $T_{ab}U_{ab}=I$. On the other hand, $U_{ab}T_{ab} \neq I$:

$$U_{ab}T_{ab}=\begin{bmatrix}0 & 0 & 0\\ 0 & 1 & 0\\ 0 & 0 & 0\end{bmatrix}$$

Thus, the special case $U_{ab}T_{ab}p=p$ will hold if and only if p does not involve reactions P_1 and P_3.

We can verify $W_{ab}R_{bp}=R_{aq}T_{ab}$ (which is crucial in the definition of R_{aq} and T_{ab}):

$$W_{ab}R_{bp}=\begin{bmatrix}0 & -1 & 0\\ 0 & 1 & 0\\ 0 & 0 & 0\end{bmatrix}=R_{aq}T_{ab}$$

It was easy for T_{ab} to have no more than one non-zero element per column, since there is only one element in each column. Similarly, in U_{ab}, there is only one element per row. The last restrictions in the above section are thus inconsequential if there is only one reaction in view A.

Procedure for Deriving Strict Abstractions

We have imposed a number of constraints on the linear transformations defining an abstraction from view B to view A. We have not addressed the question of which of these constraints are independent and what matrices or submatrices can be chosen freely. There are far too many different cases to study if one attempts to identify all the interactions and dependencies among the constraints; the cases hinge on precisely which matrices, submatrices, or matrix expressions are prespecified, and what their ranks are. Note that the abstraction-related constraints interact with the constraints imposed on each individual view, to give rise to other forms of the constraints; for example, it can be shown that the properties of reactions in view A, and the abstraction of metabolites and bioreactions from B to A, lead to the requirement $E_bV_{ab}W_{ab}R_b=0$.

There is, nevertheless, a straightforward procedure for constructing an abstraction level, such that all the constraints are satisfied. The procedure assumes that the

prior specifications are the matrices E_b, R_{bp}, and W_{ab}; it assumes that the constraints $E_bR_{bp}=0$ and $\text{rank}(W_{ab})=n_a$ are satisfied. Furthermore, the procedure assumes that the nullspace of W_{ab} is contained in the nullspace of E_b; under this assumption, the view A is more detailed than element balances, and the family of b's that map to the same $a \in A$ (by $a=W_{ab}b$) has a unique elemental composition, i.e., if $W_{ab}b=W_{ab}b'$ then $E_bb=E_bb'$.

The procedure provides one consistent solution for the remaining matrices, without regard to other possible solutions. The sequence of matrix transformations is given here, without proof.

- Compute S_{bp} and V_{ab} as pseudoinverses, $S_{bp}=(R_{bp})^+$ and $V_{ab}=(W_{ab})^+$; compute $E_a=E_bV_{ab}$
- Compute the auxiliary matrix $W_{ab}R_{bp}$; set up T_{ab} with n_p columns (its number of rows, n_q will be determined shortly); set up R_{aq} with n_a rows (its number of columns, n_q, to be determined). For each column j of $W_{ab}R_{bp}$:
  - ◊ If the column j is not proportional to any single column in R_{aq}, then add it as the next column in R_{aq}; add a row to T_{ab}, with the entry of 1 at column j and entries 0 everywhere else.
  - ◊ If the column j is equal to a multiple (with a proportionality factor K) of column i of R_{aq}, then enter that factor K as the (i,j) element of T_{ab}. Do not increase the size of R_{aq} or T_{ab}.
- Compute $U_{ab}=(T_{ab})^+$ and $S_{aq}=(R_{aq})^+$.

We reiterate that this procedure identifies just one possible choice of abstraction. The way this choice is made gives rise to additional properties that may not hold generally. For example, our extra assumption that the nullspace of W_{ab} is contained in the nullspace of E_b will make the additional relation $E_b=E_aW_{ab}$ valid for the resulting abstraction.

Example 1, concluded

We in fact chose the matrices throughout Example 1 following the above procedure. In order to illustrate the constraints, the matrices were presented in the order in which they were needed rather than the order in which they were derived.

Example 2

The abstraction route chosen in Example 1 is by no means the only way to do away with the complexes B_4 and B_5 in the above example. We could have chosen a different view A, with the same metabolites as before, but with the abstraction matrices:

$$W_{ab}=\begin{bmatrix}1&0&0&0&0\\0&1&0&0&0\\0&0&1&0&0\end{bmatrix},\quad V_{ab}=\begin{bmatrix}1&0&0\\0&1&0\\0&0&1\\0&0&0\\0&0&0\end{bmatrix}$$

These matrices simply throw away the complexes B_4 and B_5. This abstraction satisfies all the necessary restrictions, with the appropriate choice of R_{aq}.

$$R_{aq}=\begin{bmatrix}-1&0\\0&1\\-1&1\end{bmatrix}$$

The reactions would become:

Q_1: $A_1+A_3 \rightarrow 0$

Q_2: $0 \rightarrow A_2+A_3$

The transformation of substrate A_1 to product A_2 occurs when both reactions take place; $q=[1\ 1]^T$ gives $a=[-1\ 1\ 0]^T$. There is nothing fundamentally wrong with this view.

Which of the two (or other possible) modes do we usually employ? The choice would be guided in part by practical considerations, such as measurements of metabolites: If the metabolites in view A are measured such that metabolites bound into complexes (B_4 and B_5) are not included in the resulting value, then the view that ignores the complexes (Example 2) is more relevant. If, on the other hand, the measurement of the substrate A_1 includes both free substrate B_1 and bound B_4 (and A_2 and A_3 similarly include complexes) then the view we used in Example 1 is clearly appropriate.

The central message of this paper is that in our informal abstractions we do not bother to be specific, so the above modes (and many others) are all wrapped into our informal notion. If computers are to create, or reason across, levels of abstraction, the mapping must be made explicit, and we offered in this paper a framework for accomplishing this. Already, in those practical cases where numerical measurements or calculations are involved, we unavoidably fall into one of the modes - often implicitly. Numerical consistency, representational flexibility and semantic clarity stand to benefit from this framework that makes the views and abstractions explicit.

Non-Strict Abstractions

The two levels, A and B, in our analysis had a clear dominance relation: The view B was more detailed than view A (or at least as detailed as A) in every respect. This strict inclusion is the justification for many of the properties we postulated. However, we may have views which do not form a strict hierarchy. For example, view A from Example 1 and view A from Example 2 are at a similar level of complexity; none of the two is a strict refinement of the other.

Our framework allows such views to coexist and establishes relations among them. This is accomplished by relating both views to a more detailed one (in this case, the view B of Example 1 is a suitable basis). There are fewer constraints between views that are not in strict refinement relationship to each other. For example, if A and B do not form a strict abstraction, we can no longer require $W_{ab}V_{ab}=I$; but for $a \in A$ lying in a particular subspace we can require $W_{ab}V_{ab}a=a$.

Multiple views of a system thus form a graph, in which the behavior of the system in any one view can be translated to any other view, regardless of whether the view is strictly more abstract, strictly more detailed, or neither. The relation between two views does affect the invertibility of these mappings (as well as the presence of invariant subspaces).

Discussion

This paper presented an approach for the hierarchical

representation of the metabolism. The framework was illustrated with two views of a rather simple system. We consider this framework well suited for the representation of complex metabolic systems at multiple levels of detail. Multiple-level representations will be essential as we require computer models to possess more flexibility and generality. Within the proposed framework, an intelligent system could construct the appropriate view for any given problem, through iterative refinement, abstraction, or combination of given views. For a particular task, it may be necessary to lump large portions of the metabolism into very abstract views and refine others in great detail. The proposed framework defines the necessary entities and transformations for accomplishing this.

One of the benefits of multiple levels of detail is that several important informal notions can be encoded, applied, and detected, rather than simply memorized. We are referring to concepts such as *catalyst* or *inhibitor*. If, in our Example 1, we adopted only the common view A, we would need to store, as an externally given fact, the information that A_3 is the catalyst of the transformation $A_1 \rightarrow A_2$. But with both views A and B in place, the notion of a catalyst becomes explicit. The role of A_3 is apparent in the mapping between views A and B, as follows. The reaction stoichiometry $A_1 \rightarrow A_2$ maps (through matrix V_{ab}) to $\frac{1}{2}B_1 + \frac{1}{2}B_4 \rightarrow \frac{1}{2}B_2 + \frac{1}{2}B_5$, since:

$$V_{ab}[-1\ 1\ 0]^T = [-1/2\ 1/2\ 0\ -1/2\ 1/2]^T$$

Four compounds participate in this reaction from view B. Of these, B_4 and B_5 map back (through W_{ab}) to view A as follows:

B_4: $W_{ab}[0\ 0\ 0\ 1\ 0]^T = [1\ 0\ 1]^T$, or $A_1 + A_3$

B_5: $W_{ab}[0\ 0\ 0\ 0\ 1]^T = [0\ 1\ 1]^T$, or $A_2 + A_3$

Thus, in view B, we actually have A_3 participating in the reaction. This is the essential feature of a catalyst: While it is not present in the reaction itself in view A, it is found in the corresponding transformation of view B. The very need for the concept of catalyst arises because abstraction (from B to A) hides the reactions (and the role) of a particular compound (A_3), even though it does not hide the compound itself. By sneaking to the lower level and returning (a transformation encoded by the matrix $W_{ab}V_{ab}$) we recover the hidden role.

We should emphasize that we have restricted the scope of this paper in two ways, to maintain clarity. First, we were concerned only with the portion of the metabolism that does not directly involve genetic material - DNA or mRNA. Second, we ignored all issues arising from membranes, spatial arrangements or gradients, and compartmentalization. We focused on that portion of the metabolism which involves the chemical interconversion of chemical compounds (small compounds with molecular weights of up to a few hundred, as well as proteins). This portion of the metabolism of a cell is a complex network of biochemical reactions which interconvert small or large molecular-weight compounds. The bioreactions include the core metabolism that accomplishes the production of energy and the provision of the building blocks (monomers) of biological macromolecules. Since we allow the interconversion of proteins, we are including in the scope of our discussion the activation or deactivation processes operating on the enzymes themselves.

We explicitly excluded the direct study of enzyme production through transcription and translation, and, as mentioned earlier, all spatial structures. This exclusion is motivated by the need to maintain a clear and understandable exposition. The issues that prompted this paper are present in other biological processes, and the proposals set forth are relevant for the resolution of these issues throughout the domain of biochemical processes. However, the excluded processes need more careful consideration and additional analysis.

We also restricted the discussion to closed systems, but the extension to open systems is conceptually easy. The matrices would merely involve extra rows or columns for the incoming and outgoing streams. We note that a similar extension is possible for the inclusion of thermodynamic quantities, such as the Gibbs energy which can be estimated from the structures of compounds (Mavrovouniotis, 1991) and incorporated into the abstraction process.

References

Ingraham, J.L., Maaløe, O., and Neidhardt, F.C. 1983. *Growth of the Bacterial Cell.* Sunderland, Mass.: Sinauer Associates.

Karp, P. D., and Mavrovouniotis, M. L. 1994. Representing, Analyzing, and Synthesizing Biochemical Pathways. *IEEE Expert,* accepted for publication.

L. Hunter, ed. 1993. *Artificial Intelligence and Molecular Biology,* Menlo Park, Calif.: AAAI Press.

Lehninger, A.E. *Biochemistry*, 2nd ed. Worth Publishers, New York, 1975.

Mavrovouniotis, M. L. 1991. Estimation of Standard Gibbs Energy Changes of Biotransformations. *Journal of Biological Chemistry*, 266:14440-14445.

Mavrovouniotis, M. L. 1993. Identification of Localized and Distributed Bottlenecks in Metabolic Pathways. *Proceedings of the First International Conference on Intelligent Systems in Molecular Biology (ISMB-93)*, 275-283, Menlo Park, Calif.: AAAI-Press.

Mavrovouniotis, M. L., Stephanopoulos, G., and Stephanopoulos, G. 1990. Computer-Aided Synthesis of Biochemical Pathways. *Biotechnology and Bioengineering*, 36:1119-1132.

Ochs, R. S., and Conrow, K. 1991. A Computerized Metabolic Map. *J. Chem. Inf. Comput. Sci.*, 31:132-137.

Reddy, V. N., Mavrovouniotis, M. L., and Liebman, M. N. 1993. Petri Net Representations in Metabolic Pathways. *Proceedings of the First International Conference on Intelligent Systems in Molecular Biology (ISMB-93)*, p.328-336, Menlo Park, Calif.: AAAI-Press.

Strang, G. 1988. *Linear Algebra and its Applications*, 3rd ed. Academic Press.