A Graphical Calculus

Sharon Curtis and Gavin Lowe

Oxford University Computing Laboratory, Parks Road, Oxford, OX1 3QD, UK.

Abstract. We present a graphical calculus, which allows mathematical formulae to be represented and reasoned about using a visual representation. We define how a formula may be represented by a graph, and present a number of laws for transforming graphs, and describe the effects these transformations have on the corresponding formulae. We then use these transformation laws to perform proofs. We illustrate the graphical calculus by applying it to the relational and sequential calculi. The graphical calculus makes formulae easier to understand, and so often makes the next step in a proof more obvious. Furthermore, it is more expressive, and so allows a number of proofs that cannot otherwise be undertaken in a point-free way.

1 Introduction

Traditionally, mathematical formulae are written down on a single line. For example, in the relational calculus [9], given four relations P, Q, R and S, we can write $P;Q \cap R;S$ to represent the relation that relates two elements x and y iff there exist u and v such that P relates x to u, Q relates u to y, R relates x to v, and S relates v to y:

$$x(P;Q \cap R;S)y \;\Leftrightarrow\; \exists u, v \bullet x\,P\,u \wedge u\,Q\,y \wedge x\,R\,v \wedge v\,S\,y.$$

But suppose that we also want to specify that u and v are related by T. Traditional mathematics has no way of writing down such a relation in a point-free style using only the composition and intersection operators. In other words, the language of intersection and composition is expressively incomplete.

In this paper we develop a calculus of graphs for defining and reasoning about relations. For example, we represent the relation $P;Q \cap R;S$ by the graph in figure 1. Each edge represents the relation with which it is labelled; two consecutive

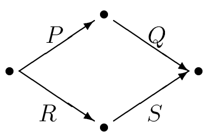

Fig. 1. A graph representing $P;Q \cap R;S$.

edges represent the composition of the corresponding relations; two paths with the same start and end points represent the intersection of the corresponding relations.

If we want to add the above condition that the intermediate points are related by T then we simply add a corresponding edge labelled T:

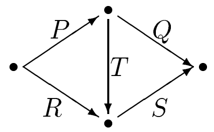

We will develop a number of *graph transformation rules*. Transforming a graph according to these rules alters the corresponding relation: for example, we will show that removing an edge from a graph makes the corresponding relation larger.

The graphical calculus provides a useful tool for doing proofs about relations: sometimes the proof without graphs is very unobvious and complicated, whereas the graphical proof is much more straightforward; and in some cases, we have proved results using the graphical calculus that we have been unable to do otherwise. The calculus gives us a way of getting at the internal structure of a relation; and because the representation is very visual, it is normally easy to see what is the correct next step in a proof.

In fact, the graphical calculus applies to more calculi than just the relational calculus. It provides a general way of representing many mathematical formulae that cannot be written down on one line in the normal way. It then provides rules for transforming these representations. We give examples of other calculi that can be represented in the graphical calculus.

In the next section we apply the graphical calculus to the relational calculus: we give a brief overview of the relational calculus, formally define how a relation can be represented by a graph, present eleven graph transformation rules, and illustrate the calculus with two examples. In section 3, we consider the sequential calculus of [10]: we describe the calculus, define how elements of the calculus can be represented by graphs, present nine graph transformation rules (eight of which are the same as in the relational calculus), and use the graph calculus to prove a result which has not otherwise been proved in the sequential calculus. In section 4 we discuss various other points of interest.

2 The relational calculus

We define a *relation* of type $A \leftrightarrow B$ to be a subset R of $A \times B$, and write xRy when $(x, y) \in R$. Composition and converse are defined in the normal way:

$$\begin{aligned} P;Q &\mathrel{\widehat{=}} \{(x,z) \mid \exists y \bullet xPy \wedge yQz\} \\ P^\circ &\mathrel{\widehat{=}} \{(y,x) \mid xPy\}. \end{aligned}$$

Union and intersection of relations are simply the corresponding set relations. We use the convention that composition binds more tightly than union and intersection. The identity relation on A is denoted by Id_A, and the universal relation on $A \times B$ by $\Pi_{A\times B}$:

$$Id_A \mathrel{\widehat{=}} \{(x,x) \mid x \in A\}$$
$$\Pi_{A\times B} \mathrel{\widehat{=}} \{(x,y) \mid x \in A \wedge y \in B\};$$

the subscripts are usually omitted, and inferred from context.

We will use two operators which return the domain and range of a relation. It is convenient to define these such that they return a relation, i.e. a set of pairs. They can be defined in a point-wise manner by:

$$dom\, R \mathrel{\widehat{=}} \{(x,x) \mid \exists y \bullet xRy\}$$
$$ran\, R \mathrel{\widehat{=}} \{(y,y) \mid \exists x \bullet xRy\}.$$

However, for calculations it is more convenient to have a point-free definition:

$$dom\, R \mathrel{\widehat{=}} Id \cap R;R^\circ$$
$$ran\, R \mathrel{\widehat{=}} Id \cap R^\circ;R.$$

We will also use the quotient operator, defined as follows:

$$R\backslash S \mathrel{\widehat{=}} \{(x,y) \mid \forall z \bullet zRx \Rightarrow zSy\}.$$

The operator may also be defined by a Galois connection:

$$R;T \subseteq S \Leftrightarrow T \subseteq R\backslash S.$$

For example, $\in\backslash\in$ represents the subset relation (where $\in$ is the set membership relation).

2.1 Representing relations by graphs

As described in the introduction, we will represent relations by graphs: each edge represents the relation with which it is labelled; composition is represented by arrows in sequence; and intersection is represented by arrows in parallel. For example, the relation $(P;Q^\circ \cap (R \cup S));T$ can be represented by:

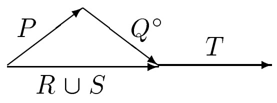

As we will see later, arrows can be reversed to give the converse of a relation, and union can be represented by splitting the graph; so the above relation may also be represented by:

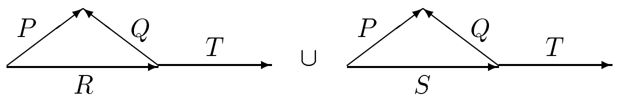

Formally, we consider graphs of the form (V, s, t, A) where V is a finite set of vertices, $s \in V$ is the source, $t \in V$ is the target, and $A \in \mathbf{P}(V \times \mathcal{S} \times V)$ is a finite set of edges labelled with elements of $\mathcal{S}$ representing relations: the edge (v, R, v') represents an edge from v to v' labelled R. When we draw a graph, the source and target will not be explicitly labelled: they will be the left-most and right-most vertices, respectively.

Note that we have no conditions concerning the connectivity of graphs. Note also that we use *sets* of edges, rather than multisets (bags); this means that a graph with two edges from v to v' labelled R is the same as the corresponding graph with only one such edge.

We can now formally define the way in which a graph represents a relation.

Definition 1. The graph $G = (\{v_0, \ldots, v_n\}, v_0, v_n, A)$ represents the relation R where

$$x\, R\, y \;\; \textit{iff} \;\; \exists x_0, \ldots, x_n \bullet x = x_0 \wedge y = x_n \wedge \forall (v_i, S, v_j) \in A \bullet x_i\, S\, x_j .$$

We call R the *interpretation* of G.

The graph represents the relation that relates x and y iff there is some way of labelling the vertices with elements such that x labels the source, y labels the target, and if there is an edge labelled S between two vertices then the corresponding elements are related by S.

For example, the graph in figure 1 relates x and y iff

$$\exists x_0, x_1, x_2, x_3 \bullet x = x_0 \wedge y = x_3 \wedge x_0\, P\, x_1 \wedge x_1\, Q\, x_3 \wedge x_0\, R\, x_2 \wedge x_2\, S\, x_3 ,$$

that is, the graph indeed represents the relation $P;Q \cap R;S$.

In the following we will use graphs when formally we mean the relations represented by those graphs. So, for example, we write $G_1 \subseteq G_2$ when the relation corresponding to G_1 is a subset of the relation corresponding to G_2; we write $G_1 \cong G_2$ when the relations are equal.

We define a number of graph transformation laws: some of these transformations leave the corresponding relation unchanged; others produce a superset of the original relation. Each of the laws may easily be proved sound with respect to the above definition.

We may enlarge the relation labelling any edge; this enlarges the relation represented by the whole graph:

Law 1 (Monotonicity) *If* $R \subseteq S$ *then*

$$(V, s, t, A \cup \{(v, R, v')\}) \subseteq (V, s, t, A \cup \{(v, S, v')\}).$$

This law allows us to incorporate techniques from the relational calculus into the graph calculus: we may use the relational calculus to prove $R \subseteq S$, and then use law 1 to replace an edge labelled R by one labelled S.

The next law uses the concept of a graph homomorphism:

Definition 2 (Homomorphism). Given graphs $G = (V, s, t, A)$ and $G' = (V', s', t', A')$, a *homomorphism* from G to G' is a function $\phi : V \to V'$ such that: (1) $\phi(s) = s'$; (2) $\phi(t) = t'$; and (3) for each edge $(v, P, v') \in A$, there is a corresponding edge $(\phi(v), P, \phi(v')) \in A'$.

For example, there is a homomorphism from the left hand graph to the right hand graph below, mapping u_0 to v_0, u_1 and u_2 to v_1, and u_3 to v_3.

Law 2 (Homomorphism) *If there is a homomorphism from G to G' then $G \supseteq G'$.*

Note that if there is a homomorphism ϕ from G to G', and another homomorphism ψ from G' to G, then $G \cong G'$. This allows us to identify the following two graphs, for example:

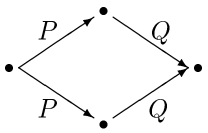

and

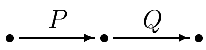

The following law states that we may always remove edges to make the corresponding relation larger. It can be proved as a corollary of the previous law, but it is sufficiently useful to be worth stating explicitly.

Law 3 (Remove edges) $(V, s, t, A \cup \{(v, R, v')\}) \subseteq (V, s, t, A)$.

An edge labelled with R may be replaced by a graph representing R:

Law 4 (Replace edge by graph) *If the relation R is represented by the graph (V', s', t', A'), and $V \cap V' = \{s', t'\}$, then*

$$(V, s, t, A \cup \{(s', R, t')\}) \cong (V \cup V', s, t, A \cup A').$$

The next four laws show how the operations of composition, intersection and union are represented in the graph calculus. If an edge is labelled by a relational composition then we may split it into two:

Law 5 (Split composition) *If v'' is a vertex not in V, then*

$$(V, s, t, A \cup \{(v, R;S, v')\}) \cong (V \cup \{v''\}, s, t, A \cup \{(v, R, v''), (v'', S, v')\}).$$

If we have two successive edges labelled with relations R and S, we may add another edge labelled with $R;S$ (this may be proved as a corollary of the previous law and the homomorphism law):

Law 6 (Composition) *If* $(v, R, v'), (v', S, v'') \in A$ *then*

$$(V, s, t, A) \cong (V, s, t, A \cup \{(v, R;S, v'')\}).$$

An edge labelled with an intersection $R \cap S$ may be replaced by two edges with the same start and end points, labelled with R and S, and vice versa:

Law 7 (Intersection)

$$(V, s, t, A \cup \{(v, R \cap S, v')\}) \cong (V, s, t, A \cup \{(v, R, v'), (v, S, v')\}).$$

If an edge of a graph is labelled with the union of two relations, R and S, then the graph may be replaced by the union of two graphs with corresponding edges labelled by R and by S:

Law 8 (Union)

$$\begin{aligned}(V, s, t, A \cup \{(v, R \cup S, v')\}) \cong\ & (V, s, t, A \cup \{(v, R, v')\}) \\ & \cup (V, s, t, A \cup \{(v, S, v')\}).\end{aligned}$$

The above laws allow a graph to be reduced to a normal form: laws 5, 7 and 8 allow compound labels to be broken down into simple labels (i.e. labels without compositions, intersections or unions); law 2 then allows redundant edges to be removed. Further, the laws—along with the observation that a graph with a single edge labelled R represents the relation R—justify our informal description of how to represent a relation by a graph.

In the above we have used graphs to represent *relations*. However, we can also use graphs to represent other sorts of mathematical formulae. Given any space $\mathcal{S}$ with operators $\cap$, $\cup$ and ; and a preorder $\subseteq$, we define a *graphical calculus* over $\mathcal{S}$ to be a calculus of graphs labelled with members of $\mathcal{S}$, such that we have some way of interpreting a graph as a member of $\mathcal{S}$, such that the above eight transformation rules are satisfied. For example, in the above we took $\mathcal{S}$ to be a space of relations, and $\cap$, $\cup$, ; and $\subseteq$ had the normal interpretations of intersection, union, relational composition and subset. In later sections we will look at other instances of graphical calculi.

Particular instances of the graphical calculus may satisfy additional laws. For example, in the relational calculus three additional laws concern the converse operator, the universal relation and the identity relation.

An edge labelled R may be reversed in direction and relabelled with the converse of R:

Law 9 (Converse)

$$(V, s, t, A \cup \{(v, R, v')\}) \cong (V, s, t, A \cup \{(v', R^{\circ}, v)\}).$$

Any two vertices are connected via the universal relation:

Law 10 (Universal relation) *If v,v' are two vertices in V, then*

$$(V, s, t, A) \cong (V, s, t, A \cup \{(v, \Pi, v')\}).$$

If two vertices are related by the identity, then they may be fused together:

Law 11 (Fusion)

$$\begin{array}{l}(V, s, t, A \cup \{(v, Id, v')\}) \cong \\ \quad (\{ren\, u \mid u \in V\}, ren\ s, ren\ t, \{(ren\, u, R, ren\, u') \mid (u, R, u') \in A\}), \\ \qquad \textit{where } ren\, u = \begin{cases} v, & \textit{if } u = v' \\ u, & \textit{otherwise.} \end{cases}\end{array}$$

The function ren renames the node v' to v.

2.2 Example: the modular law

One law commonly used in the relational calculus is the modular law, also known as the Dedekind law:

$$A \cap B;C \subseteq (A;C^\circ \cap B);C.$$

This law cannot be proved by calculation using only the properties of intersection, converse and composition in a point-free manner (there exist frameworks that satisfy these properties but not the modular law). The proof may either be executed in a point-wise fashion, or by using tabulation of relations (expressing a relation as the converse of a function followed by a function). Either method is more complicated than the proof using the graph calculus:

$A \cap B;C$

$\cong$ {Graphical representation}

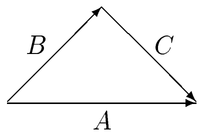

$\cong$ {Converse}

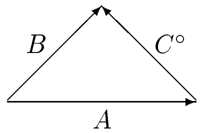

$\cong$ {Composition}

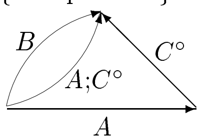

$\subseteq$ {Converse; remove edge}

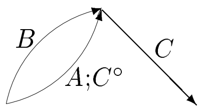

$\cong$ {Graphical representation}
$(A;C^\circ \cap B);C.$

2.3 Example: an arithmetical lemma

Another example concerns a lemma pertaining to sets of natural numbers. Writing $\ni$ for $\in^\circ$, we define the minimum with respect to a relation by:

$$min\, R = \ni \cap \in\backslash R^\circ .$$

(The pair (X, x) is in this relation if $X \ni x$ and for all y, $y \in X \Rightarrow x\, R\, y$.) Our lemma states:

$$dom \ni \; ; \in\backslash\in \; ; min \leq \;\; \subseteq \;\; min \leq \; ; \; \geq,$$

If a pair (X, y) is in the left hand side then X is non-empty, and there is some superset Y of X such that y is the minimum element of Y. The lemma states that X has a minimum x which is at least as large as y. That is, the minimum of a non-empty set of numbers is at least as large as the minimum of any superset.

Before embarking on the graphical proof, we note how the law $S\,;S\backslash R \subseteq R$ of the relational calculus may be translated into the graph calculus. If we have in our graph a subgraph

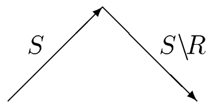

then the following arrow may be added without altering the interpretation of the graph:

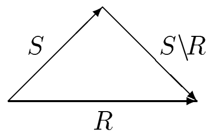

We will refer to this law as the quotient law. It may be proved using laws 1, 3 and 6.

Note also that the domains and ranges of relations may be simply represented within graphs. For example:

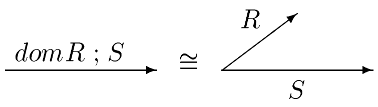

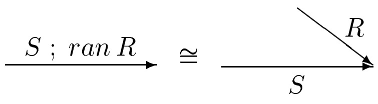

(Recall that the sources and targets of the graphs are the left-most and right-most nodes.)

The proof of the above lemma goes as follows:

$$
\begin{aligned}
& \mathit{dom}\ \ni\ ;\ \in\backslash\in\ ;\mathit{min}\ \leq \\
= & \{\text{Well-foundedness of } \geq\} \\
& \mathit{dom}(\mathit{min}\ \leq)\ ;\ \in\backslash\in\ ;\mathit{min}\ \leq \\
\cong & \{\text{Graphical representation}\}
\end{aligned}
$$

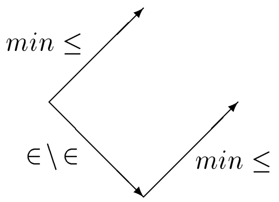

$\cong$ {Definition of min; intersection}

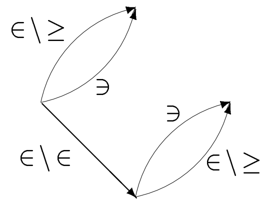

$\cong$ {Converse; quotient; converse}

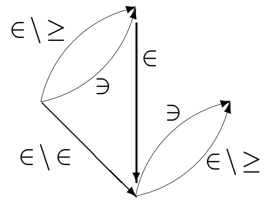

$\cong$ {Quotient}

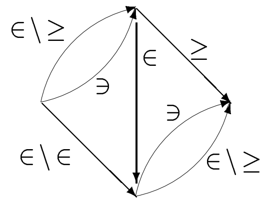

$\subseteq$ {Remove edges; definition of min}

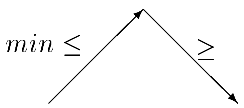

$$
\begin{aligned}
\cong & \{\text{Graphical representation}\} \\
& \mathit{min}\ \leq\ ;\ \geq .
\end{aligned}
$$

This proof illustrates a common technique in the graph calculus, namely adding all the arrows we need, and removing superfluous ones at the end.

3 Sequential calculus

The sequential calculus [10] aims to provide a common framework of algebraic laws applicable to many models of reactive systems. In this section, we examine how the sequential calculus can be modelled in the graphical calculus.

Central to the sequential calculus is the notion of an *observation*. In the calculus of intervals [1], an observation is a pair (s,t) of times—the start and termination times—with $s \leq t$. In regular expressions [8], an observation is a finite sequence of letters drawn from some alphabet A. In the regularity calculus [4], the sequences are given the structure of a group. In interval temporal logic [11], observations are functions from time intervals to states. In the traces model of CSP [6], observations are traces of visible actions. The relational calculus is also a sequential calculus, where an observation is a pair (x,y) such that x is related to y.

In each of these calculi, two observations can be composed via an associative composition operator, ";". For regular expressions, the composition operator is simply concatenation of strings. For the other calculi, composition is a partial operator; for example, in the relational calculus two observations may be composed iff the second element of the first observation is the same as the first element of the second observation; in this case the intermediate point is omitted:

$$(r,s);(s,t) = (r,t).$$

We say that observation x is a prefix of y, written $x \leq y$, if x can be extended to y:

$$x \leq y \Leftrightarrow \exists z \bullet x;z = y.$$

In each calculus, a system may be represented by a *set* of observations, termed a *sequential relation*. These form a Boolean algebra under the union and intersection operators. The composition operator may be lifted point-wise to sets:

$$P;Q \mathrel{\widehat{=}} \{p;q \mid p \in P \wedge q \in Q\}.$$

We denote the universal set of observations by Π. The main difference between the relational and sequential calculi is the lack of a converse operator in the sequential calculus.

An important concept is that of units. Each observation x has a left unit $\overleftarrow{x}$ and a right unit $\overrightarrow{x}$ such that

$$\overleftarrow{x};x \;=\; x \;=\; x;\overrightarrow{x}.$$

For example, in the relational and interval calculi, $\overleftarrow{(x,y)} = (x,x)$ and $\overrightarrow{(x,y)} = (y,y)$. We denote the set of all units by Id:

$$Id \mathrel{\widehat{=}} \{x \mid \overleftarrow{x} = x = \overrightarrow{x}\}.$$

In [10], a number of algebraic laws are developed for reasoning about sequential relations, rather than reasoning about individual observations; for example:

$$R;Id = R = Id;R, \qquad P;(Q \cup R) = P;Q \cup P;R, \qquad P;(Q \cap R) \subseteq P;Q \cap P;R.$$

3.1 Representing sequential relations by graphs

We may use the graph calculus to represent sequential relations in the obvious way. For example, if in figure 1 the labels are interpreted as sequential relations then the graph represents the sequential relation $P;Q \cap R;S$. Each edge represents the sequential relation with which it is labelled; a path through the graph represents the composition of the corresponding relations; two paths with common source and target represent the intersection of the corresponding relations.

We formalize our representation as follows:

Definition 3. The graph $G \mathrel{\widehat{=}} (\{v_0, \ldots, v_n\}, v_0, v_n, A)$ represents the sequential relation

$$\{x \mid \exists x_0, \ldots, x_n \bullet x_0 = \overleftarrow{x} \wedge x_n = x \wedge \forall i \in 0 \mathinner{..} n \bullet x_i \leq x \\ \wedge \forall (v_i, S, v_j) \in A \bullet \exists y \in S \bullet x_i;y = x_j\}.$$

We call this sequential relation the *interpretation* of G.

An observation x is in the interpretation of G if for each vertex v_i there is a corresponding observation x_i, such that:

- the observation corresponding to the source is the left unit of x;
- the observation corresponding to the target is x;
- each observation is a prefix of x;
- and for each edge (v_i, S, v_j) there is an observation y of S which when composed with x_i gives x_j.

The idea is that we start off at the source with a unit observation, and traverse the graph; on each edge we extend the observation with an observation from the edge's label, until we get to the target. Each intermediate observation is compatible with (i.e. is a prefix of) the final observation.

It is easy to prove the following theorem from the above definition:

Theorem 4 (Laws of the sequential calculus). *Each of the graph transformation laws 1–8 hold for the sequential calculus.*

The relational calculus is a particular example of a sequential calculus, so we would hope that the two ways of interpreting a graph—as a relation or as a sequential relation—are compatible; the following lemma shows that this is indeed the case.

Lemma 5. *Given a graph G labelled with relations, let R be the corresponding relation (as in definition 1) and let S be the corresponding sequential relation (as in definition 3); then:*

$$x \, R \, y \Leftrightarrow (x, y) \in S.$$

Proof. Let $G = (\{v_0, \ldots, v_n\}, v_0, v_n, A)$, and let R and S be as above. Suppose $x \, R \, y$. Then, from definition 1, there exist $y_0, \ldots, y_n$ (labelling the vertices) such that

$$x = y_0 \ \wedge \ y = y_n \ \wedge \ \forall (v_i, P, v_j) \in A \bullet y_i \, P \, y_j .$$

Then to show $(x, y) \in S$, we label the vertices with $(x, y_0), \ldots, (x, y_n)$; to meet the requirements of definition 3 we need to check:

$$\begin{aligned} &(x, y_0) = \overleftarrow{(x, y)} \ \wedge \ (x, y_n) = (x, y) \ \wedge \ \forall i \in 0 \,..\, n \bullet (x, y_i) \leq (x, y) \\ &\wedge \forall (v_i, P, v_j) \in A \bullet \exists (x', y') \in P \bullet (x, y_i);(x', y') = (x, y_j), \end{aligned}$$

which is trivial.

Conversely, suppose $(x, y) \in S$. Then, from definition 3, there exist $(x_0, y_0), \ldots, (x_n, y_n)$ such that

$$\begin{aligned} &(x_0, y_0) = \overleftarrow{(x, y)} \ \wedge \ (x_n, y_n) = (x, y) \ \wedge \ \forall i \in 0 \,..\, n \bullet (x_i, y_i) \leq (x, y) \\ &\wedge \forall (v_i, P, v_j) \in A \bullet \exists (x', y') \in P \bullet (x_i, y_i);(x', y') = (x_j, y_j), \end{aligned}$$

from which it follows that $\forall i \in 0 . . n \bullet x_i = x$. To show $x \, R \, y$, we label the vertices with $y_0, \ldots, y_n$; to meet the requirements of definition 1 we need to check:

$$x = y_0 \ \wedge \ y = y_n \ \wedge \ \forall (v_i, P, v_j) \in A \bullet y_i \, P \, y_j ,$$

which is again trivial.

3.2 Local linearity

Many sequential calculi satisfy an additional axiom, that of *local linearity*. This is expressed at the level of observations as follows:

Definition 6 (Local linearity).

$$\begin{aligned} x;y = x';y' \Rightarrow \; & \exists w \bullet x;w = x' \wedge w;y' = y \\ & \vee \exists w \bullet x';w = x \wedge w;y = y' . \end{aligned}$$

This may be expressed as a pair of commuting diagrams:

$$\begin{tikzcd} & \bullet \arrow[dd, "w"] \arrow[dr, "y"] & \\ \bullet \arrow[ur, "x"] \arrow[dr, "x'"'] & & \bullet \\ & \bullet \arrow[ur, "y'"'] & \end{tikzcd} \qquad \text{or} \qquad \begin{tikzcd} & \bullet \arrow[dr, "y"] & \\ \bullet \arrow[ur, "x"] \arrow[dr, "x'"'] & & \bullet \\ & \bullet \arrow[uu, "w"] \arrow[ur, "y'"'] & \end{tikzcd}$$

Lifting the axiom of local linearity to the level of sets of observations has proved difficult. One formulation is

$$\begin{aligned} P;Q \cap R;S = \; & (P \cap R;\Pi);Q \cap R;(\Pi;Q \cap S) \\ & \cup (P;\Pi \cap R);S \cap P;(Q \cap \Pi;S). \end{aligned}$$

However, this formulation does not seem to be strong enough for all our requirements.

In the graph calculus, the axiom of local linearity can be expressed as follows: if we have a graph G containing two edges with start points v_i and v_j, and common end point v_k,

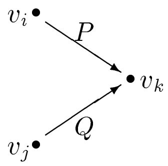

then we can add an edge labelled with the universal relation $\varPi$ either from v_i to v_j or from v_j to v_i:

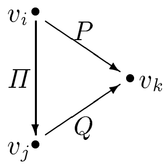

or

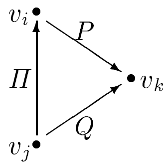

(Note that the above pictures may be subgraphs of the complete graph.) This is formalized as follows:

Law 12 (Local linearity) *If* $(v_i, P, v_k), (v_j, Q, v_k) \in A$ *then*

$$(V, s, t, A) \cong (V, s, t, A \cup \{(v_i, \varPi, v_j)\}) \cup (V, s, t, A \cup \{(v_j, \varPi, v_i)\}).$$

Proof. Obviously, from the law for removing edges (law 3), the right hand side is a subset of the left hand side.

Let $G = (V, s, t, A)$ be such that $(v_i, P, v_k), (v_j, Q, v_k) \in A$. Let R be the sequential relation represented by G, and let $x \in R$. Then there exist $x_0, \ldots, x_n$ such that

$$\begin{array}{l} x_0 = \overleftarrow{x} \wedge x_n = x \wedge \forall i \in 0\,..\,n \bullet x_i \leq x \\ \wedge \forall (v_i, S, v_j) \in A \bullet \exists y \in S \bullet x_i;y = x_j. \end{array}$$

In particular, there exist $p \in P$ and $q \in Q$ such that $x_i;p = x_j;q = x_k$. Then by the axiom of local linearity,

$$\exists w \bullet x_i;w = x_j \wedge w;q = p \qquad \text{or} \qquad \exists w \bullet x_j;w = x_i \wedge w;p = q.$$

In the first case, consider the graph $G' = (V, s, t, A \cup \{(v_i, \varPi, v_j)\})$. Then x is a member of the sequential relation represented by G': to show this we simply label G' with the same observations $x_0, \ldots, x_n$ as above, and note that $\exists w \in \varPi \bullet x_i;w = x_j$.

Similarly, in the second case x is a member of the sequential relation represented by $(V, s, t, A \cup \{(v_j, \varPi, v_i)\})$.

3.3 Example: the 3-$\Diamond$ law

We will now use the above graph transformation rule to prove a law known as the 3-$\Diamond$ law. Define

$$\Diamond X \mathrel{\widehat{=}} \varPi;X;\varPi.$$

Note that $\Diamond X$ corresponds to the "somewhere X" of interval temporal logic: it contains all observations that include an element of X as a subobservation.

The 3-$\Diamond$ law states:

$$P;Q;R \cap \Diamond X \subseteq P;(Q;R \cap \Diamond X) \cup (P;Q \cap \Diamond X);R \cup \Diamond(X \cap \Diamond Q).$$

That is, if an observation of X occurs sometime during an observation of $P;Q;R$, then either it occurs during $Q;R$, or it occurs during $P;Q$, or Q occurs during X. Much effort has gone into proving this law using the standard axioms of the sequential calculus, but without success.

Using the graph calculus version of the axiom of local linearity, the proof is extremely straightforward:

$P;Q;R \cap \Diamond X$

$\cong$ {Graph representation}

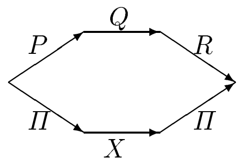

$\cong$ {Local linearity}

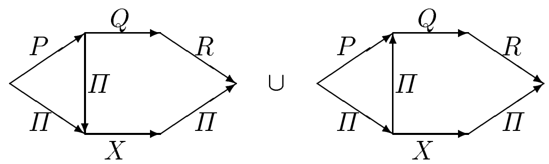

$\subseteq$ {Local linearity applied to second graph}

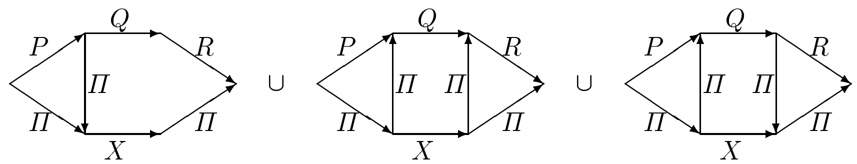

$\subseteq$ {Removing edges}

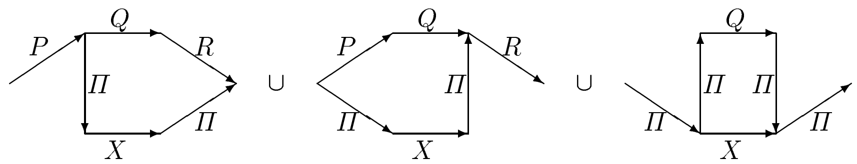

$\cong$ {Relations corresponding to graphs}

$P;(Q;R \cap \Diamond X) \ \cup \ (P;Q \cap \Diamond X);R \ \cup \ \Diamond(X \cap \Diamond Q).$

4 Discussion

In this paper we have presented a graphical calculus. We have described how to represent mathematical formulae—for example relations or sequential relations—by graphs. We have presented rules for transforming graphs and explained how these rules affect the corresponding formulae. In this final section we discuss a few other points of interest.

4.1 Related work

Brown and Hutton [2] have developed a calculus of pictures, oriented towards circuit design. Their pictures are built up from basic cells and wires using sequential composition, intersection and reciprocation. They give a semantics to pictures in terms of relations, in a manner very similar to our approach. In [2, 3] it is shown that their calculus is complete in that two pictures are equivalent with respect to their transformation rules if and only if they represent the same relation for all interpretations of the basic cells; this proof proceeds by viewing pictures as arrows in a unitary pretabular allegory [5].

Our approach is more general: their approach is restricted to calculi with intersection, composition and converse, whereas ours includes the union operator, or can exclude the converse operator. Furthermore, their approach is more oriented towards treating basic cells as simply symbols, and proving circuits equivalent in an automated manner [7]; whereas our calculi—particularly the relational calculus—are more oriented towards using the properties of the basic relations themselves in order to manually prove results concerning those relations. The Brown–Hutton pictures seem to be the easier to use for circuit design, whereas our graphs are suitable for more abstract calculi.

4.2 Other graphical calculi

We believe that many other calculi can also be fitted into the framework of the graphical calculus. For example, consider graphs labelled with positive numbers—to represent lengths—and where the interpretation of a graph is the length of the shortest path from source to target. This is a graphical calculus when one interprets the operations $\cup$, $\cap$, ; and $\subseteq$ as maximum ($\sqcup$), minimum ($\sqcap$), addition ($+$) and less-than ($\leq$), respectively. We leave it to the reader to check that the graph transformation laws 1–8 are satisfied.

While this calculus is not very interesting in its own right, it does provide some evidence that the graph calculus may be of more general applicability.

We have tried to provide a general framework for others to produce their own graphical calculus: they have only to formally define the way in which a graph represents a formula in their setting, check that the eight graph transformation laws 1–8 hold, and derive other laws particular to their calculus. Any law in the underlying calculus will have a counterpart in the graphical calculus (because of the monotonicity law), but in some cases the graphical law will be stronger (for example, the local linearity law of the sequential calculus).

4.3 Advantages of the graph calculus

One major advantage of the graph calculus is that expressive power is increased, allowing us to define and reason about more formulae. For example, Tarski [9] gives an example of a predicate not expressible as a sentence of the relational calculus:

$$\exists w, x, y, z \bullet x\, R\, y \wedge x\, R\, z \wedge x\, R\, w \wedge y\, R\, z \wedge y\, R\, w \wedge w\, R\, z.$$

We may express this predicate in the graphical calculus as follows:

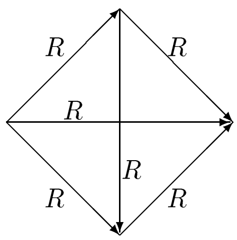

$$\not\cong \quad \emptyset.$$

If the relation represented by the above graph is non-empty then the vertices in the graph can be labelled by w, x, y and z (clockwise from the bottom) such that $x\, R\, y \wedge x\, R\, z \wedge x\, R\, w \wedge y\, R\, z \wedge y\, R\, w \wedge w\, R\, z$, and conversely. We are grateful to C. A. R. Hoare for referring us to this example.

The extra expressive power of the graphical calculus makes some proofs possible that cannot be done otherwise, for example the proofs of the modular law and the 3-$\diamond$ law above. Even in short proofs, the steps taken often result in intermediate graphs that are not directly translatable back to the underlying calculus. Even when the extra expressive power of the graphical calculus is not used, graphical proofs can be easier because they give a very visual representation of formulae, and this can make the next step more obvious.

Some formulae themselves may be simpler as graphs. For example, in the relational calculus, formulae involving *dom*, *ran*, Id or Π are often greatly simplified in the graphical representation.

Products of relations are also easily represented, by graphically interpreting their definition in terms of projections:

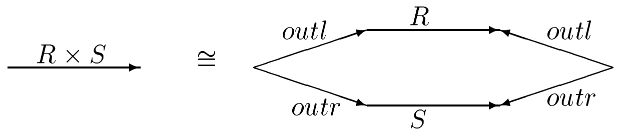

where *outl* and *outr* are the normal projection relations. This yields the pictorially intuitive idea of products being represented as parallel arrows; the graphical representation makes it easier to reason about each element of the pair separately.

4.4 Generalizing the graph calculus

In this paper we defined a graphical calculus to be defined over any structure $\mathcal{S}$ with operators $\cap$, $\cup$ and ; and a preorder $\subseteq$ such that the laws 1–8 hold. The question then arises as to whether we need all these laws, or even whether we need more. It may be that we can find calculi that we would like to consider as graphical calculi, but which satisfy only some of these laws; for example, if the underlying structure includes only composition and intersection operators, then we would not have the union law. Further investigations will be needed to find the laws common to all graphical calculi.

So far we have been considering graphs with two special vertices, the source and the target. We can easily generalise this to allow graphs with k special nodes, representing a k-ary relation. Tarski [9] gave another example of a predicate not expressible in the relational calculus:

$$\forall x, y, z \bullet \exists u \bullet x\, R\, u \wedge y\, R\, u \wedge z\, R\, u.$$

This can be represented using a graph representing a ternary relation, with the three outermost nodes representing the three components of the relation:

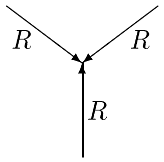

$$\cong \qquad \Pi,$$

The graph represents the ternary relation that relates x, y and z if there is some way of labelling the central node with u such that $x\, R\, u \wedge y\, R\, u \wedge z\, R\, u$; if this is the universal relation then R satisfies the above predicate.

The graph calculus can be extended to use *hyperedges* within graphs (i.e. edges with more than two ends) to represent k-ary relations; then composition of relations would be simply represented by hyperedges sharing nodes. We hope that a suitable pictorial representation could be found that would make reasoning about such relations easier.

Acknowledgements

We are grateful to Tony Hoare, He Jifeng, Oege de Moor, Graham Hutton, and the anonymous referees for useful comments, suggestions and encouragement.

References

1. Stephen Brien. A time-interval calculus. In R. Bird, C. Morgan, and J. Woodcock, editors, *Mathematics of Program Construction*, volume 669 of *LNCS*. Springer Verlag, 1992.
2. Carolyn Brown and Graham Hutton. Categories, allegories and circuit design. In *Ninth Annual IEEE Symposium on Logic in Computer Science*, pages 372–381, 1994.
3. Carolyn Brown and Alan Jeffrey. Allegories of circuits. In *Third International Symposium, Logical Foundations of Computer Science*, volume 813 of *Lecture Notes in Computer Science*, pages 56–68. Springer-Verlag, 1994.

4. E. W. Dijkstra. The unification of three calculi. In M. Broy, editor, *Program Design Calculi*, pages 197–231. Springer Verlag, 1993.
5. Peter Freyd and Andre Scedrov. *Categories, Allegories*. North-Holland, 1990.
6. C. A. R. Hoare. *Communicating Sequential Processes*. Prentice Hall, 1985.
7. Graham Hutton, Erik Meijer, and Ed Voermans. A tool for relational programmers. Available on WWW from `http://www.cs.ruu.nl/people/graham/allegories.txt`, 1994.
8. S. Kleene. Representation of events in nerve nets and finite automata. In Shannon and McCarthy, editors, *Automata Studies*, pages 3–42. Princeton University Press, 1956.
9. A. Tarski. On the calculus of relations. *Journal of Symbolic Logic*, 6:73–89, 1941.
10. Burghard v. Karger and C. A. R. Hoare. Sequential calculus. *Information Processing Letters*, 53:123–130, 1995.
11. Chaochen Zhou, C. A. R. Hoare, and Anders P. Ravn. A calculus of durations. *Information Processing Letters*, 40:269–276, 1992.

This article was processed using the LaTeX macro package with LLNCS style