# Pandora: A Reasoning Toolbox using Natural Deduction Style

KRYSIA BRODA, JIEFEI MA, GABRIELLE SINNADURAI and ALEXANDER SUMMERS, *Department of Computing, Imperial College London, South Kensington Campus, Exhibition Road, London SW7 2AZ, UK. E-mails: {kb,jm103,apgs,ajs300m}@doc.ic.ac.uk*

Abstract

Pandora is a tool for supporting the learning of first order natural deduction. It includes a help window, an interactive context sensitive tutorial known as the ''e-tutor'' and facilities to save, reload and export to LaTeX. Every attempt to apply a natural deduction rule is met with either success or a helpful error message, providing the student with instant feedback. Detailed electronic logs of student usage are recorded for evaluation purposes. This paper describes the basic functionality, the e-tutor, our experiences of using the tool in teaching and our future plans.

Keywords: natural deduction, Fitch box proof, first order logic, predicate logic, teaching, learning, e-learning, reasoning about programs
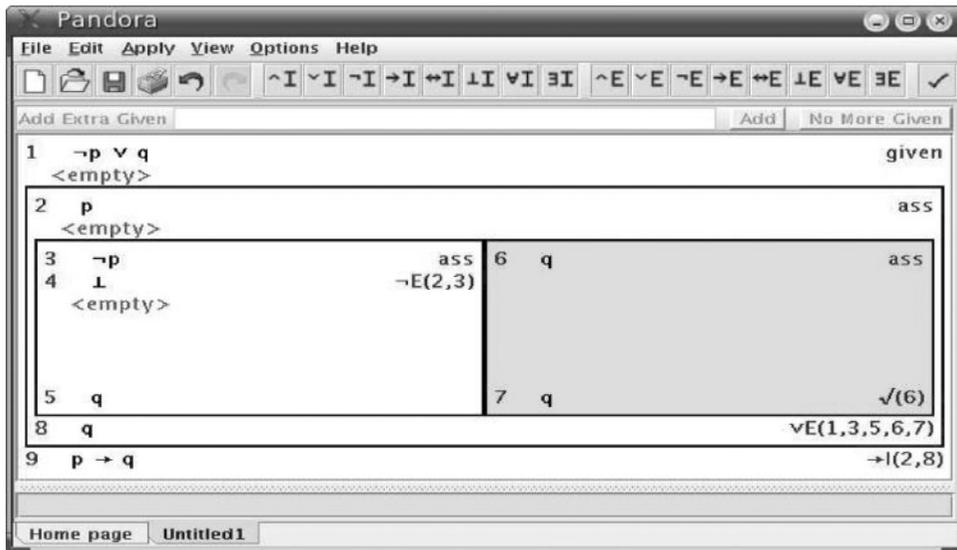
1 History and Context

Natural deduction has been taught to first year undergraduates in the Department of Computing at Imperial College since 1991. Most students find the high degree of rigour required in formal natural deduction proofs daunting. Pandora (Proof Assistant for Natural Deduction using Organised Rectangular Areas) is a learning support tool designed to guide

the construction of natural deduction proofs [1]. The acronym was coined by Lloyd Kamara. It was conceived in 1996 and is based on the Fitch-style "proof box notation" for natural deduction as described in [2], which has been the recommended text for the course since it was developed from the course lecture notes. An advantage of this notation is that the boxes enforce similar scoping rules to block structured imperative programming languages. The first version of Pandora was written by a final year student, Patrick Chan, using Tcl/Tk in 1996 and was available for students to use in tutorials but was not promoted in lectures. A second version was written in an early release of Java in 1999 by another project student, Dan Ehrlich. This second version was more robust and was used for demonstrations in class as well as in tutorials. It was available for students to download onto their own machines and promoted vigorously and was consequently widely used. The third and current version, also written in Java, was developed as a group project in 2002 and has since been further enhanced by several summer students, Jiefei Ma, Alexander Summers and Tom Weedon. It provides the basic functionality together with a help module, a tutorial which provides the user with dynamically generated advice, and the facility to export proofs as LATEX. An applet version is available for use via the Pandora web site [8] and we hope to make it available for download in the near future.

2   Basic Functionality

Pandora provides learning support to guide students in their construction of a natural deduction proof of a conclusion or goal from given premisses. It allows the user to reason "forwards", that is, from one or more given formulas to deduce another formula using one of the rules, and to reason "backwards", that is, to reduce one of the current goals to one or more subgoals from which the current goal can be deduced using one of the rules. The rules are all the usual introduction and elimination rules of first order natural deduction with equality, plus a few derived rules found useful by more advanced users. The full set can be found in [2].

We will explain the usual use of the tool by working through a small example to derive $p \rightarrow q$ from the premiss $\neg p \lor q$. On starting a new proof Pandora first requires the premisses and goal to be given. These are typed by the user into a text box. In our example there is only one premiss so we type ~p|q then click on the "No More Given" button. Pandora then requires us to give the goal. We type p>q into the text box. Pandora checks that these are well-formed formulas and then displays the initial state of the partial proof with the premisses at the top and goal at the bottom. This is shown on the left below.

$$
\begin{array}{ll}
1 & \neg p \lor q \qquad \text{given} \\
  & <empty> \\
2 & p \rightarrow q \qquad \text{conclusion}
\end{array}
\qquad\qquad
\begin{array}{ll}
1 & \neg p \lor q \qquad\quad \text{given} \\
  & <empty> \\
2 & p \qquad\qquad\quad \text{ass} \\
  & <empty> \\
3 & q \\
4 & p \rightarrow q \qquad \rightarrow\!\mathcal{I}(2{,}3)
\end{array}
$$

Notice that we input the symbols for "not", "or" and "implies" as the plain text characters ~, | and > but Pandora displays them as ¬, ∨ and →. We experimented with using buttons to input the symbols but found that plain text was more convenient and having to learn a second set of symbols for the connectives did not present any problem to our students. Now for the proof. First we apply the → introduction rule backwards by selecting the goal, $p \rightarrow q$, and clicking on the → $\mathcal{I}$ button.

Pandora displays the new proof state, shown on the right above. The box shows the scope of the assumption, p, which is discharged by the $\rightarrow \mathcal{I}$ rule and the conclusion, $p \rightarrow q$, now has the justification $\rightarrow \mathcal{I}(2,3)$.

Next we would like to apply the $\vee$ elimination rule to split the premiss $\neg p \vee q$ into its two cases. The current goal is q and we want to maintain this so we must use the rule backwards by selecting the goal, q, and clicking on the $\vee \mathcal{E}$ button then clicking on the premiss $\neg p \vee q$. The new proof state is displayed.

$$
\begin{array}{llll}
1 & \neg p \vee q & & \text{given} \\
& <empty> & & \\
\hline
2 & p & & \text{ass} \\
& <empty> & & \\
\hline
3 & \neg p \quad\text{ass} & 5 & q \quad\text{ass} \\
& <empty> & & <empty> \\
4 & q & 6 & q \\
\hline
7 & q & & \vee\mathcal{E}(1,3,4,5,6) \\
8 & p \rightarrow q & & \rightarrow\mathcal{I}(2,7)
\end{array}
$$

In the right hand box we have q as an assumption and q as the goal so we can just use the "tick" rule to note that the goal is proven. So we select the goal q, then click on the ✓ button, then on the assumption q. Pandora displays the new proof state in which the right hand box is greyed out and has the $<empty>$ line removed to indicate that that part of the proof is now completed.

To save space we will not show it yet, but first apply a rule in the left hand box. This will be our first application of a rule forwards. We have both p and $\neg p$ in scope so we can deduce $\bot$ by the $\neg$ elimination rule. To do this in Pandora we select the $<empty>$ line in the left box, then click the $\neg\mathcal{E}$ button, then click on the two formulas which give the contradiction. The new proof state displayed by Pandora is the one in the screen shot on the first page.

Next we apply $\bot$ elimination backwards from the goal q in the left hand box by selecting the goal then clicking on the $\bot\mathcal{E}$ button. This gives $\bot$ as the new goal. As usual, Pandora would display the new proof state.

Finally we use the tick rule again to note that we have $\bot$ and we need to show $\bot$. The proof is now complete. Pandora removes all the empty lines and greys out the proof and is intelligent enough to remove the last application of the tick rule which is not needed in the final proof. Note that the line numbers and references to them in the justifications were consistently updated as the proof emerged. Note also that in the completed proof every line has a justification.

$$
\begin{array}{llll}
1 & \neg p \vee q & & \text{given} \\
\hline
2 & p & & \text{ass} \\
\hline
3 & \neg p \quad\text{ass} & 6 & q \quad\text{ass} \\
4 & \bot \quad \neg\mathcal{E}(2,3) & & \\
5 & q \quad \bot\mathcal{E}(4) & 7 & q \quad \checkmark(6) \\
\hline
8 & q & & \vee\mathcal{E}(1,3,5,6,7) \\
9 & p \rightarrow q & & \rightarrow\mathcal{I}(2,8)
\end{array}
$$

If at any stage the student tries to apply a rule wrongly they are given an error message. Pandora's basic functionality as outlined above is augmented by a help window,

an interactive tutorial known as the "e-tutor" and an export to LaTeX facility. The student can save and reload partially constructed or complete proofs and work on several proofs at once, each in its own proof window. They can undo proof steps with the undo button allowing them to try alternative proof strategies. All the proof states shown above were generated using the export to LaTeX option.

3   The e-tutor

On loading Pandora, the user is presented with the Pandora home page, which offers various options, including loading a previously stored proof, starting a new proof or starting an e-tutorial. When the e-tutorial is started, four tutorials of propositional exercises are available. The first consists of fixed exercises and is useful in teaching when we demonstrate the tool in class or run a laboratory session in which all the students are required to attempt the same exercise. The other three, known as "easy", "medium" and "hard", consist of sets of five exercises randomly selected at run time.

In normal mode, Pandora just responds to the student's actions by either applying the rule they request or giving appropriate error messages. The student's strategy is not checked and a student who applies a rule correctly but unwisely, turning an achievable goal into an unprovable goal, is not warned. When an exercise is attempted in the e-tutor the level of support is greatly enhanced. Hints and explanations are provided, a warning and counter-model are given if the student applies a rule which creates an unprovable goal, and the friendly e-tutor will do a step itself if asked.

We will explain the use of the e-tutor by following a hypothetical student's attempt to derive Peirce's law $((p \rightarrow q) \rightarrow p) \rightarrow p$ which is one of the fixed tutorial exercises. In the e-tutor the view is split into a left window in which the proof is constructed and a right window in which advice is offered. Initially the proof window just contains the goal $((p \rightarrow q) \rightarrow p) \rightarrow p$ and the advice window just contains the remark "I am always here to help you" and a clickable link "I need some advice". Suppose our student clicks the link. The tutor suggests "Arrow introduction rule" and the suggestion is itself a clickable link to an explanation of the rule. On clicking this link the student is given an explanation of the usual use of the rule and what will happen if it is applied in this particular case and another clickable link "How do I do it?". On clicking this link the student is told exactly what to click on to apply the rule and given two more clickable links. In this case the advice is to select the goal, $((p \rightarrow q) \rightarrow p) \rightarrow p$, and click on the $\rightarrow \mathcal{I}$ button. The links are "Do the steps for me!" and "Show the advice again". When our lazy student selects the former the proof window is updated as shown in the proof state below and the advice window goes back to showing the remark "I am always here to help you" and a clickable link "I need some advice".

$$
\begin{array}{ll}
& < empty > \\
\hline
1 \quad (p \rightarrow q) \rightarrow p & \text{ass} \\
& < empty > \\
2 \quad p & \\
\hline
3 \quad ((p \rightarrow q) \rightarrow p) \rightarrow p & \rightarrow\mathcal{I}(1,2)
\end{array}
$$

Our student can again ask for advice and this time will be offered the choice of "Not Not Elimination", "Law of Excluded Middle" or "Proof by Contradiction". Again these are

clickable links to explanations of the rules and on selecting them further offers of links "How do I do it?" and "Do the steps for me!" would be supplied as above. But suppose our overconfident student reasons that the goal p could be proved by $\rightarrow\mathcal{E}$ using $(p \rightarrow q) \rightarrow p$ if $p \rightarrow q$ can be proved and so decides to ignore the advice and instead applies the $\rightarrow\mathcal{E}$ rule backwards by clicking on the goal p then the $\rightarrow\mathcal{E}$ button then the assumption $(p \rightarrow q) \rightarrow p$. The friendly e-tutor allows the rule to be applied in the proof window but responds with a warning in the advice window "Be careful, one or more of your goals are not provable!" and offers the links "Why?", "Undo last step" and "I want to carry on anyway". If the student asks why, the tutor provides a counter-model as shown in the screenshot below. If the student decides to carry on anyway the tutor reports that "No advice is available at the moment because the goal is not provable".
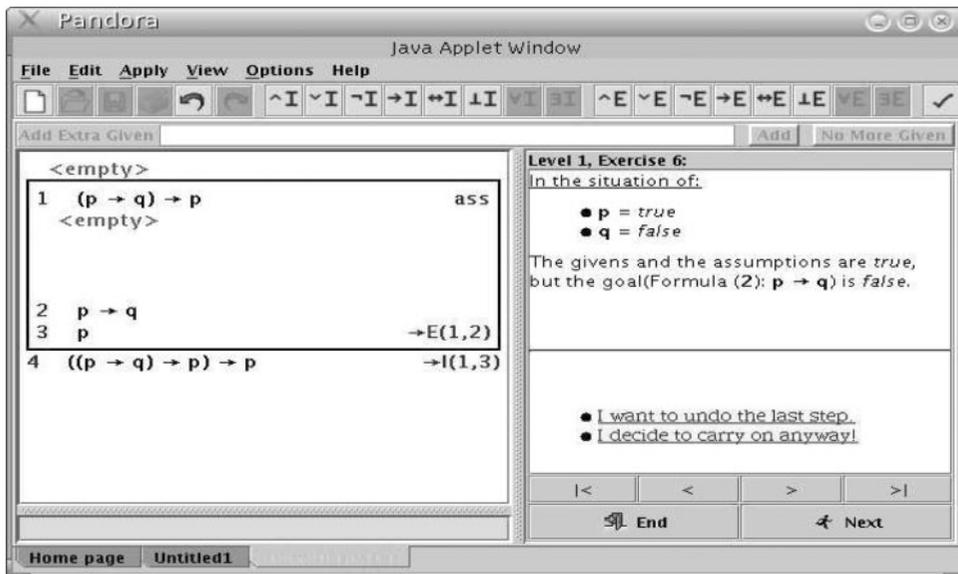
Our student can undo the incorrect step and apply instead one of the advised rules, and then continue either with or without the help of the tutor. The e-tutor just sits quietly presenting the link "I need some advice" unless the student either clicks that link or makes a move which results in an unprovable goal.

All the e-tutor's advice is dynamically generated and context sensitive. Whenever asked for advice it selects the rules to advise by finding all the applicable rules and filtering out those whose application would result in an unprovable goal. It recommends Not Not Elimination, Proof by Contradiction (PC) and Excluded Middle only as a last resort. Since only propositional exercises are available in the e-tutor, provability of goals is decidable and counter-models are easily generated.

4   Using Pandora in Teaching

We have two cohorts of students learning logic in their first year. One class (JMC) is small (20-25 students) and studying for a Joint Mathematics and Computing degree, whereas the other class (COMP) is large (over 100 students) and studying for a Computing degree. Pandora has been used differently in teaching the two cohorts.

The JMC group is taught by the third author who uses Pandora in an integrated way to teach natural deduction which constitutes about 25% of the logic module. Initially, the propositional natural deduction rules are presented and hand written examples of proofs are given. Only after they have seen several proofs and tried a few on paper themselves are students introduced to the tool. Pandora is demonstrated using several of the same examples so the students can focus on how to drive the user interface rather than on how to prove the theorems. Over the next few weeks, in the lab and tutorials, we give the students many exercises, some assessed, and introduce them to the first order and equality rules. The course finishes with a "driving test" consisting of ten problems which the students have one hour to attempt under exam conditions. Last year the test problems were:

1. $\vdash (p \to (q \to r)) \to ((p \to q) \to (p \to r))$
2. $\neg p \vee q \vdash (p \to q)$
3. $\neg(p \to q) \vdash p \wedge \neg q$
4. $p \to (q \to (r \vee s)), \neg(r \vee t), (s \wedge q) \to t \vdash q \to \neg p$
5. $\forall x \neg p(x) \vdash \neg \exists x p(x)$
6. $\neg \exists x p(x) \vdash \forall x \neg p(x)$
7. $\neg \forall x p(x) \vdash \exists x \neg p(x)$
8. $\neg \exists x \exists y (\neg x = y) \vdash \forall x \forall y (p(x, y) \to p(y, x))$
9. $\forall x (x = a \vee x = b), \neg(a = b), g(a) = b, \forall x \forall y (g(x) = g(y) \to x = y) \vdash g(b) = a$
10. $\forall x \forall y \forall z (r(x, y) \wedge r(y, z) \to r(x, z)), \forall x (r(x, a) \vee r(x, b)), r(a, b) \vdash \exists y \forall x r(x, y)$

Pandora has been enhanced with various features to facilitate automatic marking. For assessed coursework we provide the students with exercise "skeletons". These include the given premises and conclusion together with a "magic number" which encodes any restrictions we care to impose on the rules they are allowed to use. For example, we can disable certain derived rules (e.g. the PC rule) or all the predicate rules, which may be a distraction when they are proving a propositional assertion. The magic number also encodes the current state of the proof so the students cannot take the number from one proof and put it on another. The students download the skeletons from our web-based Continuous Assessment Tracking System and the first time they save a proof (usually when they have completed it!) their identity is also coded into the magic number. Whenever the proof is saved subsequently, the coded identity is unchanged so the first saver can be checked against the submitter of the proof. We are pleased to say that to date we have never found a mismatch between saver and submitter even though the students are not aware that we can make this check. The electronically gathered proofs from both the driving test and assessed coursework are checked for correctness and converted to LaTeX by "text to text" command line programs included in the Pandora package. We can thus produce a report for each student and a summary of results for their tutors with minimal human intervention.

The experience given to the larger class (COMP) has varied over time. Initially use of the current version of Pandora was not encouraged, with the result that few students tried it. Later there was light encouragement and a demo, and more students tried it out, some using it to do their coursework. Only recently has Pandora been seriously encouraged, but there is still no driving test and no requirement to use it for coursework.

For both cohorts, all exercises done throughout the term, including the driving test for the JMC, are essentially formative. The main summative assessment is an end of year written examination.

5   Evaluation

We have put considerable effort into evaluating Pandora and used three methods.

Firstly, we asked the students what they think of it, both verbally and using anonymous feedback forms. The feedback is generally encouraging and students say and write that they enjoy using Pandora and find it useful. The JMC classes, who learned Pandora thoroughly, enjoyed using Pandora more and gave more encouraging feedback than the COMP cohort. Moreover, the feedback from the COMP cohort improved as Pandora was more actively encouraged. We concluded that perhaps students tell us what they think we want to hear! Students also made useful criticisms: for example the first release of the current version only had the facility to undo the last rule application but the facility to repeatedly undo steps was added by popular demand.

Secondly, we compared performance on the written exam by the two cohorts. This did not give the clear-cut advantage to Pandora users that we hoped for and in fact there was little difference in terms of marks between the two cohorts. We believe the reasons for this were that (i) natural deduction, being only 25% of the course, did not form a large part of the exam and (ii) the questions had to be small enough to be completed under exam conditions. We believe that a difference between the cohorts may have shown up if the questions had been more substantial. We noticed a difference in style between Pandora users, who were much more at home with using rules backwards and did not make "arbitrary" assumptions which they had no hope of discharging, and non-users, who mainly reasoned forwards and frequently made arbitrary assumptions. We had feared that Pandora users may find it hard to adapt to writing proofs by hand but it turned out that the users were more precise syntactically in their hand written proofs than were the non-users.

For the third evaluation method we electronically recorded detailed logs of the students' use of Pandora; essentially we recorded every "click" they make so that we can see in detail how they actually used it. This allowed us to see which rules were applied, whether undo, help or the e-tutor were used, and so on. We computed proportions of attempted rule applications that were correct and counted error rates for various types of errors. Some of the results from analysing the logs came as a disappointment in that they showed that the help and the e-tutor's facilities were used rather infrequently. The logs also showed a surprisingly high failure rate in students' attempts to apply the rules. A small number of students had virtually no failures but many had almost as many failures as successes. Analysis of the logs showed that many students were not selecting the $< empty >$ line or a goal line before applying a rule. Comparing the logs for the driving test with those for the previous work we were pleased to observe that, with experience, the proportion of failed rule applications decreased. The logs yielded detailed information about the common errors made by students for each rule but, perhaps surprisingly, not many general problems could be diagnosed. The one fact that really jumped out of the data was that students frequently tried to apply the $\neg\mathcal{I}$ rule backwards to a formula which was not a negation, whereas it was comparatively rare for them to try applying $\rightarrow\mathcal{I}$ backwards to a formula

which was not an implication. We believe they were confusing the $\neg\mathcal{I}$ rule with the derived PC rule.

Overall the evaluation has taught us that Pandora is well liked and considered useful by the students, but that to improve the learning outcomes we should modify our teaching in a number of ways. We have run the course again since the evaluation and made some improvements; notably (i) we explained how to avoid what the logs showed to be the common pitfalls in applying the rules, (ii) when demonstrating the individual rules we emphasised that either the $< empty >$ or $< goal >$ line needs to be selected **before** clicking the rule application button, and (iii) we gave an advertisement for the e-tutor and built-in help as part of the initial demonstration. The demo of the e-tutor producing the counter-model initially elicited surprise from the students and they were very keen for me to select "I want to carry on anyway". When the e-tutor said "No advice is available at the moment because the goal is not provable" the students laughed because they could identify with its predicament. We think the laugh indicated a warm feeling towards the e-tutor as well as an understanding of the semantic checks it was making.

6   Future direction 1: Pandora and Prolog

We teach Prolog [4] in the third term of the first year, by which time our students are very familiar with natural deduction and Pandora. If natural deduction is applied to the subset of first order logic consisting only of Horn clauses, it is quite easy to simulate the action of a Prolog interpreter. So why not use Pandora to help their understanding of backtracking and clausal reasoning?

Prolog can be introduced without resolution by viewing Prolog derivations of query Q from program P as natural deduction proofs of the goal Q from premises P. A Prolog program consists of definite clauses with the general form $A \leftarrow B_1 \wedge \ldots \wedge B_n$, where A, $B_1, \ldots, B_n$ are atoms. In the case $n = 0$ the clause has the form A and is called a fact; otherwise it is called a rule. Any variables in a clause are implicitly universally quantified. A query is a conjunction of atoms, in which variables are implicitly existentially quantified. Prolog finds all solutions for an existential query by back-tracking. This is conveniently modelled in Pandora by use of undo.

Currently in $\forall\mathcal{E}$ forwards and $\exists\mathcal{I}$ backwards only a ground term can be substituted for the bound variable. These rules will be generalised to allow the introduction of a fresh unknown in place of the bound variable. The tick rule currently just checks that two ground formulas match. For Prolog this must be extended to include unification between atoms. Similarly backwards application of the rule $\rightarrow\mathcal{E}$ must allow for the atomic goal to unify with the consequent atom of the major ($\leftarrow$) premise. For convenience backwards application of $\wedge\mathcal{I}$ will be generalised for an arbitrary number of conjuncts.

Below is an illustration of the deduction of a query from a five clause program, in which it is assumed that a derivation has already been found (using b for x), then undone, and that a second derivation is sought.

When a rule is applied which binds an unknown the replacement of variables by values must happen throughout the proof. This overwriting means that the unknowns are lost from the proof state. The diagram below is not itself a proof state because we have left some of the unknowns in so that a shadow of earlier states can be seen. One such

replacement, which was caused by the `tick` rule used to justify line 7, is applied to line 10 and indicated by $\Uparrow$.

1	$arc(a,b)$		program fact
2	$arc(a,d)$		program fact
3	$arc(b,c)$		program fact
4	$\forall u \forall v[path(u,v) \leftarrow arc(u,v)]$		program rule
5	$\forall u \forall v \forall w[path(u,w) \leftarrow arc(u,v) \wedge path(v,w)]$		program rule
6	$path(?u1,?w1) \leftarrow arc(?u1,?v1) \wedge path(?v1,?w1)$		$\forall \mathcal{E}(5)$

		8	$path(?u2,?v2) \leftarrow arc(?u2,?v2)$	$\forall \mathcal{E}(4)$
		9	$arc(b,c)$	$\surd$ (3)
		10	$path(b,c)$ (if $?u2/b,?v2/c$)	$\rightarrow\mathcal{E}(8,9)$
			$\Uparrow$ caused by justification (7)	
7	$arc(?x1,?v1)$ $\surd$(1) if $?x1/a,?v1/b$	10	$path(?v1,c)$	

11	$arc(?x1,?v1) \wedge path(?v1,c)$	$\wedge\mathcal{I}(7,10)$
12	$path(?x1,c)$ (if $?u1/?x1,?w1/c$)	$\rightarrow\mathcal{E}(6,11)$
13	$\exists x[path(x,c)]$	$\exists\mathcal{I}(12)$

7   Future direction 2: Reasoning about Programs

We are also working on enhancing Pandora with Hoare logic for reasoning about programs. A clear simple account of Hoare logic with its semantic justification was given by Winskel [14] and formalised in the generic theorem prover Isabelle [6] by Nipkow et al [7, 10]. The Hoare rules as used in the Isabelle theory are:

skip $: \vdash \{P\}SKIP\{P\}$
assg $: \vdash \{\lambda s.P(s[x \mapsto as])\}x = a\{P\}$
semi $: [\![\vdash \{P\}c\{Q\}; \vdash \{Q\}d\{R\}]\!] \Longrightarrow \vdash \{P\}c;\ d\{R\}$
if $: [\![\vdash \{\lambda s.Ps\&bs\}c\{Q\}; \vdash \{\lambda s.Ps\&\neg bs\}d\{Q\}]\!] \Longrightarrow \vdash \{P\}IF\ b\ THEN\ c\ ELSE\ d\{Q\}$
while $: \vdash \{\lambda s.Ps\&bs\}c\{P\} \Longrightarrow \vdash \{P\}WHILE\ b\ DO\ c\{\lambda s.Ps\&\neg bs\}$

where s varies over states, that is, functions from memory locations to values, and arithmetic expressions and boolean expressions are represented as functions from states to natural numbers and states to booleans respectively. We have done some work on adapting the formalism to Pandora style but have not yet started the java implementation of the upgraded tool.

1	$sx = a \wedge sy = b$	pre
	$< empty >$	
P1	`x = x + y; y = x-y; x = x - y`	
2	$sx = b \wedge sy = a$	$< goal >$

We present a partial correctness proof of a simple swap program to give an indication of what we are trying to do. We start with the precondition, the program and the $< goal >$ or hoped for postcondition as shown above. We assume that ; associates to the left so that the

outer connective is the second. We can apply the rule `semi` by clicking on the program line `P1` then the rule `semi` to get the proof state below.

1	$sx = a \land sy = b$	pre
	$< empty >$	
2	$?P$	$< goal >$
P1	$\text{x} = \text{x} + \text{y}; \text{y} = \text{x} - \text{y}; \text{x} = \text{x} - \text{y}$	

> | 3 | $?P$ | ass |
> | | $< empty >$ | |
> | P2 | $\text{x} = \text{x} + \text{y}; \text{y} = \text{x} - \text{y}$ | |
> |
> | 4 | $?MID$ | $< goal >$ |
> | 5 | $?MID$ | ass |
> | | $< empty >$ | |
> | P3 | $\text{x} = \text{x} - \text{y}$ | |
> |
> | 6 | $?Q$ | $< goal >$ |

7	$?Q$	semi P1 (3,4)(5,6)
	$< empty >$	
8	$sx = b \land sy = a$	$< goal >$

Here the $?P$, $?Q$ and $?MID$ are unknown formulas which can be unified with any formula using the `tick` rule. At this point rule `semi` could be applied again to line `P2` or we can use the `tick` rule to unify the formula at line 8 with $?Q$ and remove the redundant line. If we do the latter we can then use the `assg` rule on `P3` to bring the $< goal >$ at line 6 above the code line `P3` to get the proof state below.

1	$sx = a \land sy = b$	pre
	$< empty >$	
2	$?P$	$< goal >$
P1	$\text{x} = \text{x} + \text{y}; \text{y} = \text{x} - \text{y}; \text{x} = \text{x} - \text{y}$	

> | 3 | $?P$ | ass |
> | | $< empty >$ | |
> | P2 | $\text{x} = \text{x} + \text{y}; \text{y} = \text{x} - \text{y}$ | |
> |
> | 4 | $?MID$ | $< goal >$ |
> | 5 | $?MID$ | ass |
> | | $< empty >$ | |
> | 6 | $sx - sy = b \land sy = a$ | $< goal >$ |
> | P3 | $\text{x} = \text{x} - \text{y}$ | |
> | 7 | $sx = b \land sy = a$ | assg P3 (6) |
> | 8 | $sx = b \land sy = a$ | semi P1 (3,4)(5,7) |

Now we use the `tick` rule to unify the formula at line 6 with $?MID$. This causes the redundant line to be removed and the bottom box which is now complete to be greyed out. Next we must use the `semi` rule on `P2` and then we can proceed as above using `tick` to unify

and `assg` to bring the goal above the code until we get up to the precondition and complete the proof. The resulting proof is shown below.

1	$sx = a \wedge sy = b$	pre
2	$sy = b \wedge sx = a$	TML(1)
P1	$\mathtt{x = x + y;\, y = x - y;\, x = x - y}$	
3	$sy = b \wedge sx = a$	ass
P2	$\mathtt{x = x + y;\, y = x - y}$	
4	$sy = b \wedge sx = a$	ass
5	$sy = b \wedge (sx + sy) - sy = a$	TMA(4)
P3	$\mathtt{x = x + y}$	
6	$sy = b \wedge (sx - sy) = a$	assg P3 (5)
7	$sy = b \wedge (sx - sy) = a$	ass
8	$sx - (sx - sy) = b \wedge (sx - sy) = a$	TMA(7)
P4	$\mathtt{y = x - y}$	
9	$sx - sy = b \wedge sy = a$	assg P4 (8)
10	$sx - sy = b \wedge sy = a$	semi P2 (4,6)(7,9)
11	$sx - sy = b \wedge sy = a$	ass
P5	$\mathtt{x = x - y}$	
12	$sx = b \wedge sy = a$	assg P5 (11)
13	$sx = b \wedge sy = a$	semi P1 (3,10)(11,12)

Note that when the rule `assg` was applied backwards at line 9 over program line `P4` and backwards at line 6 over program line `P3` the resulting expressions at lines 8 and 5 needed simplifying to get the lines 7 and 4 respectively. In the long term we hope to be able to add an arithmetic module to Pandora to enable it to simplify, or at least check the correctness of simplifications, itself. In order to get started we will initially build in a "trust me arithmetic" (TMA) rule. Applications of this rule will be accepted by the tool if some simple semantic plausibility checks are passed. Note also that in going backwards to the precondition from line 2, which like line 3 was obtained by unification with line 4, we use the `TML` rule which stands for "trust me logic" and will be used for simple logical checking.

For these ideas to scale up for use with larger programs we need the ability to collapse and expand parts of the proof so they take up less screen space when they are not being actively worked on. Also, undo/redo needs to be made smarter so that a small change required in a loop invariant does not cause all the work since the invariant was first given to be lost. This is a useful challenge - simple enough to provide hope of progress but substantial enough to give us some confidence that good solutions will scale up to larger proof systems.

The big advantage of Pandora over text-based interactive theorem provers such as Isabelle is that the student can see the proof emerge graphically and non-essential details like line numbers are automatically dealt with. Unlike the final proof produced by Isabelle, which can be viewed as an "assembly code" program for generating a proof rather than a proof itself, the Pandora proof is human readable. Its overall structure can be seen at a glance and the detailed steps can easily be followed.

8 Related Work

Pandora is not the only teaching tool for first order natural deduction using Fitch-style boxes. Two others of which we are aware are Jape [5] and CPT-AProS [9].

Jape was originally developed about the same time as Pandora by Richard Bornat and Bernard Suffrin. It has recently been enhanced and although it does not currently have an

e-tutor to give hints it does have the useful ability to look for first order counter-models. It includes some unification on formulas, for example applying the $\vee\mathcal{I}$ rule to the formula A gives $A \vee ?B$ and the unknown can later be unified with any formula. Jape silently carries out the `tick` rule whenever it can, even when unification is involved. From a teaching point of view we believe this is a draw-back as the students are not in control of the whole proof construction. Also, when this happens the proof state ''jumps'' which is visually confusing.

This paper is an extension of [3] presented at Salamanca where the Carnegie Mellon Proof Tutor Project (CPT) was also presented by Wilfried Sieg [12]. CPT provides a Proof Lab in which students construct natural deduction proofs as part of a web-based course on modern formal logic including other interactive learning environments and video clips. We understand it is currently being enhanced with an e-tutor to give dynamically generated strategic advice using an efficient algorithm for finding natural deduction proofs developed as part of the Automated Proof Search (AProS) project. The algorithm is based on the intercalation calculi [13] and is complete for both propositional and predicate logic.

References

[1] K. Broda, S. Eisenbach and L. Kamara. Tools for Natural Deduction, *Proceedings of First Australian Conference on Computer Science Education*, ACSE96, CACM, pp 119-126, 1996.

[2] K. Broda, S. Eisenbach, H. Khoshnevisan, S. Vickers. *Reasoned Programming*, Prentice-Hall, 1994. (Available at [8].)

[3] K. Broda, J. Ma, G. Sinnadurai and A. Summers. Pandora: Natural Deduction made Easy, in [11], pp 11-14.

[4] W. Clocksin and C. Mellish. *Programming in Prolog: Using the ISO Standard*, Springer-Verlag, 2003.

[5] http://jape.comlab.ox.ac.uk:8080/jape/

[6] http://www.cl.cam.ac.uk/Research/HVG/Isabelle/

[7] http://www.cl.cam.ac.uk/research/hvg/Isabelle/dist/library/HOL/IMP/index.html

[8] http://www.doc.ic.ac.uk/pandora/

[9] http://www.phil.cmu.edu/projects/apros/

[10] T. Nipkow. Winskel is (almost) Right: Towards a Mechanized Semantics Textbook, *Foundations of Software Technology and Theoretical Computer Science*, ed. V. Chandru and V. Vinay, Springer LNCS 1180, pp 180-192, 1996.

[11] *Proceedings of the second international congress on tools for teaching logic*, Eds: Maria Manzano, Belen Perez Lancho and Ana Gil, Salamanca, 2006.

[12] W. Sieg. STRATEGIC THINKING: web-based and dynamically tutored, in [11], pp 199-200.

[13] W. Sieg and J. Byrnes. Normal Natural Deduction Proofs (in classical logic), *Studia Logica* 60: pp 67-106, 1998.

[14] G. Winskel. *The Formal Semantics of Programming Languages*, MIT, 1993.