

# The Fox Project: Advanced Development of Systems Software

R&D Status Report  
April 1 to June 30, 1997

School of Computer Science  
Carnegie Mellon University  
Pittsburgh, PA 15213

This research is sponsored by the Defense Advanced Research Projects Agency, DoD, through ARPA Order 8313, and monitored by ESD/AVS under contract F19628-95-C-0050. Views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the United States Government.

The long-term objectives of the Carnegie Mellon Fox Project are to improve the design and construction of systems software and to further the development of advanced programming language technology. We use principles and techniques from the mathematical foundations of programming languages, including semantics, type theory, and logic, to design and implement systems software, including operating systems, network protocols, and distributed systems. Much of the implementation work is conducted in the Standard ML (SML) language, a modern functional programming language that provides polymorphism, first-class functions, exception handling, garbage collection, a parameterized module system, static typing, and a formal semantics. This Project involves several faculty members and spans a wide range of research areas, from (1) experimental development of systems software to (2) advanced compiler development to (3) language design.

**DISTRIBUTION STATEMENT A**

Approved for public release;  
Distribution Unlimited

**DTIC QUALITY INSPECTED 1**

19970703 154

# 1 Research Progress

For each of the three areas listed above, we report on the research accomplishments during the second calendar quarter of 1997, and the research objectives for the third quarter of 1997.

## 1.1 Experimental Development of Systems Software

### Accomplishments (April-June):

- Completed the implementation of the ATM FoxNet. This is one of the major deliverables for the Fox Project.
- Improved the performance of the FoxNet, to achieve a total TCP/IP throughput of 7.3Mb/s on an Ethernet and 23Mb/s on ATM, with the FoxNet running in user space. (In both cases, TCP checksums are disabled.)
- Updated the source code for the FoxNet so that it conforms to the new SML'97 language specification. In particular, the source now obeys the so-called "value restriction" on polymorphic values. This will facilitate the use of the TIL compiler (currently under development) to compile the FoxNet system.

### Objectives (July-September):

- Conduct a thorough performance evaluation of the ATM FoxNet, and begin tuning its performance. This will provide a basis for analyzing performance bottlenecks, as well as language and compiler deficiencies due to the Standard ML language.

## 1.2 Language Design

### Accomplishments (April-June):

- Completed a working implementation of "refinement type analysis" for the Standard ML language. Refinement types hold out the promise of allowing more detailed typechecking and programmer specifications in the ML language, thereby improving reliability and composability of systems software.
- Made a complete re-implementation of proof-reconstruction for the full LF language. The LF language is used to encode proofs in Proof-Carrying Code, and hence the efficiency of proof-reconstruction is of great importance. The re-implementation will allow further optimizations than was possible before.
- George Necula presented a thesis proposal entitled "Compiling with Proofs." This work will test the feasibility of the use of TIL-style compiler technology to build Certifying Compilers that automatically generate Proof-Carrying Code binaries.

- Chris Okasaki has signed a contract with Cambridge University Press to publish a book on functional data structures and algorithms. This book will be based largely on Okasaki's thesis work, and is likely to be the standard text on the subject.

### **Objectives (July-September):**

- Design and make a first implementation of optimizations for LF type-checking and proof reconstruction.
- Build and demonstrate a working prototype of a certifying compiler for a simple programming language.
- Complete revisions of the book on functional data structures for Cambridge University Press.

### **Noteworthy publications:**

- "Catenable Double-Ended Queues," by Chris Okasaki, ACM SIGPLAN International Conference on Functional Programming, June 1997.
- "Implementing Bit-addressing with Specialization," by Scott Draves, ACM SIGPLAN International Conference on Functional Programming, June 1997.
- Chris Okasaki is an invited lecturer at NSF-CISE workshop on Integrating Recent Research Results into the Undergraduate Curricula, July 8, 1997.

## **1.3 SML Compiler and System Development**

### **Accomplishments (April-June):**

- Completed the front-end of the TIL compiler, and tested its correctness by compiling several simple programs, using last year's prototype code generator to generate machine code.
- Completed the typechecker for the main intermediate language (called MIL) of the TIL compiler. This provides an important internal consistency check for the TIL compiler. It also facilitates exchange of compiler components with a research group at Cornell University, who are also engaged now in joint TIL compiler development. In particular, the Cornell group has agreed to implement the main MIL-level optimizations for the TIL compiler.
- Adapted the Bell Labs ML-RISC code generation system so that it would accommodate the TIL compiler. This code generator is now working in limited test cases.

### **Objectives (July-September):**

- Continue development of the TIL front-end, so that it accepts the entire SML'97 language.

- Complete the adaptation of the ML-RISC code generator, and connect it up with the TIL front-end..

### **Noteworthy publications:**

- Robert Harper co-chaired a workshop on "Types In Compilation." This workshop was extremely successful, with over 80 attendees. This is a strong indication of the influence of the TIL compiler design.

## **2 Capital Equipment Purchases**

- 1 Cisco Access Server 2511, \$1,257.00
- 1 264020U IBM ThinkPad Model 560 P133, 8MB, 2.1GB, 12.1TFT, Win95, \$4,345.00
- 2 ST15150N 2 4GB 3.5" Narrow SCSI-2 Disks, x2C13520 in 3.5" Dual Bay Enclosure, \$3,850.00
- 1 ST15150N 1 4GB 3.5" Narrow SCSI-2 Disks, C13520 in 3.5" Dual Bay Enclosure, \$1,030.00
- 2 HP LaserJet 5M Printer with Postscript Level II, \$2,944.00

## **3 Key Personnel Changes**

- None.

## **4 Noteworthy Meetings**

- IFIP WG2.8 meeting (Harrogate, England, June 2-6). Robert Harper and Chris Okasaki were invited participants in the meeting on Functional Programming.
- IEEE Security and Privacy conference (Oakland, CA, May 4-7). George Necula participated in a panel entitled "Security in New Innovative Operating Systems" where he presented the "Proof-Carrying Code" technique. He also participated in an OS Security "Birds of a Feather" session organized by DARPA.
- Fourth ACM Computer and Communications Security conference (Zurich, Switzerland, April). Peter Lee participated in a panel on "Programming Language Techniques for Security", where he presented the "Proof-Carrying Code technique.

## 5 Administrative Data

```
<ADMINISTRATIVE_AND_FINANCIAL_DATA>
<DATE_PREPARED>
27 JUN 1997
</DATE_PREPARED>
<AWARD_NUMBER>
F19628-95-C-0050
</AWARD_NUMBER>
<AWARD_AGENT>
Hanscom Air Force Base
</AWARD_AGENT>
<AWARD_TITLE>
The Fox Project: Advanced Languages for Systems Software
</AWARD_TITLE>
<ACTUAL_START_DATE>
08 JUN 1995
</ACTUAL_START_DATE>
<ACTUAL_START_DATE_COMMENT>
</ACTUAL_START_DATE_COMMENT>
<OFFICIAL_AWARD_END_DATE>
30 JUN 1998
</OFFICIAL_AWARD_END_DATE>
<CURRENT_AWARD_PROFILE_BY_FISCAL_YEAR>
<FY>95</FY><BASE>922,250</BASE>
<FY>96</FY><BASE></BASE>
<OPT_NAME>Option1</OPT_NAME><OPT_AMT>964,201</OPT_AMT>
<FY>97</FY><BASE></BASE>
<OPT_NAME>Option1</OPT_NAME><OPT_AMT></OPT_AMT>
<OPT_NAME>Option2</OPT_NAME><OPT_AMT>1,008,341</OPT_AMT>
<OPT_NAME>Option4</OPT_NAME><OPT_AMT>275,005</OPT_AMT>
<FY>98</FY><BASE></BASE>
<OPT_NAME>Option1</OPT_NAME><OPT_AMT></OPT_AMT>
<OPT_NAME>Option2</OPT_NAME><OPT_AMT></OPT_AMT>
<OPT_NAME>Option4</OPT_NAME><OPT_AMT></OPT_AMT>

</CURRENT_AWARD_PROFILE_BY_FISCAL_YEAR>
<OPTIONS_UNFUNDED_BY_FISCAL_YEAR>
<FY>98</FY><OPT_NAME>Option
  4</OPT_NAME><OPT_AMT></OPT_AMT>
</OPTIONS_UNFUNDED_BY_FISCAL_YEAR>

<TOTAL_FUNDS_PROVIDED_TO_DATE>
2,400,497
</TOTAL_FUNDS_PROVIDED_TO_DATE>
<ACTUAL_FUNDS_EXPENDED_TO_DATE>
1,795,000
</ACTUAL_FUNDS_EXPENDED_TO_DATE>
<ACTUAL_FUNDS_UNEXPENDED_TO_DATE>
605,497
</ACTUAL_FUNDS_UNEXPENDED_TO_DATE>
```

<DATE\_CURRENT\_FUNDING\_EXPENDED>  
30 NOV 1997  
</DATE\_CURRENT\_FUNDING\_EXPENDED>

<FUNDING\_FOR\_NEXT\_FY>  
769,300  
</FUNDING\_FOR\_NEXT\_FY>  
<DATE\_OF\_FINANCIAL\_DATA>  
30 JUN 1997  
</DATE\_OF\_FINANCIAL\_DATA>

<ANYTHING\_ELSE\_YOU\_NEED>

</ANYTHING\_ELSE\_YOU\_NEED>

</ADMINISTRATIVE\_AND\_FINANCIAL\_DATA>