

Emerald Heterogeneous Migration

April 9, 2001

Motivation: why migrate?

- o fault tolerance? restart task in a new place, typically from checkpoint
- o load balancing? move objects off of overloaded nodes
- o performance? reduce communication via migration
idea: replace lots of RPCs with one expensive migration (possibly two, to move it back)
- o main alternative: careful placement of new objects to balance load
- o overall: advantages are not that clear and there is great deal of complexity...

Easy migration:

- o use only machine independent form, such as Java byte codes. Interpret or JIT compile threads as needed -- this seems like a viable solution given JIT compilation to address performance
- o interpreted languages (eg. tcl/perl) are relatively easy to migrate, but still need to package up all of the state. (why easier in tcl than perl? tcl uses only strings)

(homogeneous) native code migration:

- o has performance win: threads run at full speed
- o big challenge: how to package up state
- o code and function pointers is easy in the homogeneous case, since code base is the same
- o hard part: registers, stack, objects
- o challenge: which values are scalars (e.g. ints) and which are pointers? Pointers must be converted to equivalent values on new host. implies: 1) must know for every register, stack frame location, and heap location whether something is a pointer (at the time of migration!), and 2) must be able to tell where the corresponding object is...
- o pointers into local objects? either convert to global references or make a stub object on the new host and point to it.
- o unions are very hard. Emerald disallows unions that are not all scalar or all pointers, otherwise can't tell if you need to change the pointer. Unlike conservative GC, must have a 100% answer (or don't migrate the object)
- o indirect addressing (e.g. via stack pointer) is helpful and needs no translation
- o compiler produces template for every object and every stack frame to indicate the pointers
- o also need a convention for registers... easiest one is to partition registers into pointers and scalars and don't mix them. more complex: have stack template define register use as well

heterogeneous native code migration:

- o hard problem 1: code pointers and function pointers now need translation
- o basic solution: map code pointers to machine independent locations then map them back to machine dependent on the new architecture
- o hard problem 2: no all code locations HAVE a viable translation. Examples: reordered code, code that is hoisted on one platform but not on another, atomic instructions that correspond to many instructions on another platform means that you can't map from the middle of the sequence back to the host with the atomic instruction
- o general solution: define bus stops: points at which code maps 1-1. for example, at bus stop, both versions have executed the same semantic set of instructions.
- o Key insight: between bus stops, the programs need not correspond 1-1 -- compilers are free to optimize as they like.
- o Typical bus stop locations: procedure/system call boundaries, the bottom of loops
- o required compiler support: bus stop mappings, use of register, temp variables, etc.

Bridging code (not implemented):

- o idea: migrate at points with no direct mapping
- o solution: 1) define location in UNoptimized code, 2) bridge to that point, 3) migrate, 4) bridge from there to optimized location on target host
- o not implemented or evaluated

other tidbits:

- o when do you need a network format? only if the two sides differ and you don't want $O(n^2)$ conversions (for n architectures)
- o code is immutable, which means it need not be moved
- o how to gain control upon next method call? reset stack pointer so that stack allocation fails (check has to be there anyway) Note: can't really use page faults or modify the code; in theory could modify the return address of the active procedure, but it may not return for a long time, and there is a race condition
- o must track the USE of variables/registers at every bus stop: different types may require different conversions (not true in homogeneous case)
- o need to make sure different compilers produce the same bus stops (in the same order) -- this was done by hand
- o struct conversion is by recursive descent (much like RMI marshalling)
- o very hard to debug!! errors show up only during migration and may be path or machine dependent; also very hard to test every bus stop
- o stack frames are translated at the receiver youngest to oldest, which is the opposite of how they need to be allocated; not clear why this is required. They thus have to move frames after everything is translated. In theory, they could avoid the movement using two phases, one for sizing (quick) and one for translation...
- o could you migrate C threads? probably not... unless you use a type-safe subset -- otherwise you can't tell how to translate something (not even homogeneous migration)

- o how to transfer objects that use local resources (such as open TCP connections or the CD-ROM)? probably best to use a proxy object on the remote host
- o migration in Java in the presence of JIT compilation? (easy without JIT, since JVM is machine independent) need to migrate only at byte code boundaries, which may require bridging code (to get to the next one), or at least identification of which points map to such boundaries
- o de-optimization is also very useful for debugging