

Shared Memory Based JITLink Memory Manager

Student: Anubhab Ghosh

Mentors: Vassil Vassilev, Lang Hames, Stefan Gränitz

JITLink

- JITLink is a Just-In-Time linker.
 - It takes multiple object code units and links them together.
 - It constructs the result directly in memory.
 - The resulting code is usually immediately run.
- It uses a LinkGraph as memory representation.
 - It consists of nodes like Addressable, Block, Symbol.
 - Relocations are represented by edges.
 - Sections consists of symbols and blocks.
- It can work with two different processes.
 - An executor process that is running the resultant code.
 - A controller process that performs the linking and controls the executor.
 - The communication happens through Executor Process Control, an RPC scheme.

Memory management

- Memory allocation is performed using the JITLinkMemoryManager interface.
- It has 3 steps
 - Allocate
 - Reserves address space
 - Finalize
 - Copies link result from working memory to executor
 - Runs initialisation actions
 - Deallocate
 - Runs deinitialization actions
 - Deallocate memory
- Initialization and deinitialization actions are just functions that will be executed in the context of the target process.

```
class JITLinkMemoryManager {  
public:  
    class FinalizedAlloc {  
        orc::ExecutorAddr release();  
    };  
  
    class InFlightAlloc {  
        virtual void abandon(OnAbandonedFunction OnAbandoned) = 0;  
        virtual void finalize(OnFinalizedFunction OnFinalized) = 0;  
    }  
  
    virtual void allocate(const JITLinkDylib *JD, LinkGraph &G,  
                          OnAllocatedFunction OnAllocated) = 0;  
  
    virtual void deallocate(std::vector<FinalizedAlloc> Allocs,  
                           OnDeallocatedFunction OnDeallocated) = 0;  
};
```

When multiple processes are involved, this is implemented with the **EPCGenericJITLinkMemoryManager** and **SimpleExecutorMemoryManager**.

The Executor Process side

```
31 // Simple page-based allocator.
32 class SimpleExecutorMemoryManager : public ExecutorBootstrapService {
33 public:
34     virtual ~SimpleExecutorMemoryManager();
35
36     Expected<ExecutorAddr> allocate(uint64_t Size);
37     Error finalize(tpctypes::FinalizeRequest &FR);
38     Error deallocate(const std::vector<ExecutorAddr> &Bases);
```

```
73 struct SegFinalizeRequest {
74     WireProtectionFlags Prot;
75     ExecutorAddr Addr;
76     uint64_t Size;
77     ArrayRef<char> Content;
78 };
79
80 struct FinalizeRequest {
81     std::vector<SegFinalizeRequest> Segments;
82     shared::AllocActions Actions;
83 };
```

```
24 Expected<ExecutorAddr> SimpleExecutorMemoryManager::allocate(uint64_t Size) {
25     std::error_code EC;
26     auto MB = sys::Memory::allocateMappedMemory(
27         Size, nullptr, sys::Memory::MF_READ | sys::Memory::MF_WRITE, EC);
33     return ExecutorAddr::fromPtr(MB.base());
34 }
```

```
36 Error SimpleExecutorMemoryManager::finalize(tpctypes::FinalizeRequest &FR) {
109     // Copy content and apply permissions.
110     for (auto &Seg : FR.Segments) {
128         char *Mem = Seg.Addr.toPtr<char *>();
129         memcpy(Mem, Seg.Content.data(), Seg.Content.size());
130         memset(Mem + Seg.Content.size(), 0, Seg.Size - Seg.Content.size());
131         assert(Seg.Size <= std::numeric_limits<size_t>::max());
132         if (auto EC = sys::Memory::protectMappedMemory(
133             {Mem, static_cast<size_t>(Seg.Size)},
134             tpctypes::fromWireProtectionFlags(Seg.Prot)))
135             return BailOut(errorCodeToError(EC));
136         if (Seg.Prot & tpctypes::WPF_Exec)
137             sys::Memory::InvalidateInstructionCache(Mem, Seg.Size);
138     }
```

- Implemented using a bootstrap service.
- 3 primary functions: allocate, finalize and deallocate
- Deallocation actions are also transferred during finalization.

The Controller Process

The controller process side is implemented in EPCTGenericJITLinkMemoryManager.

It mainly consists of RPC calls to the methods of SimpleExecutorMemoryManager.

```
void finalize(OnFinalizedFunction OnFinalize) override {
    tpctypes::FinalizeRequest FR;
    for (auto &KV : Segs) {
        assert(KV.second.ContentSize ≤ std::numeric_limits<size_t>::max());
        FR.Segments.push_back(tpctypes::SegFinalizeRequest{
            tpctypes::toWireProtectionFlags(
                toSysMemoryProtectionFlags(KV.first.getMemProt())),
            KV.second.Addr,
            alignTo(KV.second.ContentSize + KV.second.ZeroFillSize,
                    Parent.EPC.getPageSize()),
            {KV.second.WorkingMem, static_cast<size_t>(KV.second.ContentSize)}});
    }

    // Transfer allocation actions.
    std::swap(FR.Actions, G.allocActions());

    Parent.EPC.callSPSWrapperAsync<
        rt::SPSSimpleExecutorMemoryManagerFinalizeSignature>(
        Parent.SAs.Finalize,
        [OnFinalize = std::move(OnFinalize), AllocAddr = this->AllocAddr](
            Error SerializationErr, Error FinalizeErr) mutable {
            // FIXME: Release abandoned alloc.
            if (SerializationErr) {
                cantFail(std::move(FinalizeErr));
                OnFinalize(std::move(SerializationErr));
            } else if (FinalizeErr)
                OnFinalize(std::move(FinalizeErr));
            else
                OnFinalize(FinalizedAlloc(AllocAddr));
        },
        Parent.SAs.Allocator, std::move(FR));
}
```

EPC implementation under the hood

```
836 int FromExecutor[2];
837
838 pid_t ChildPID;
839
840 // Create pipes to/from the executor..
841 if (pipe(ToExecutor) != 0 || pipe(FromExecutor) != 0)
842     return make_error<StringError>("Unable to create pipe for executor",
843                                     invertibleErrorCode());
844
845 ChildPID = fork();
846
847 if (ChildPID == 0) {
848     // In the child...
849
850     // Close the parent ends of the pipes
851     close(ToExecutor[WriteEnd]);
852     close(FromExecutor[ReadEnd]);
853
854     // Execute the child process.
855     std::unique_ptr<char[]> ExecutorPath, FDSpecifier;
856     {
857         ExecutorPath = std::make_unique<char[]>(OutOfProcessExecutor.size() + 1);
858         strcpy(ExecutorPath.get(), OutOfProcessExecutor.data());
859
860         std::string FDSpecifierStr("filedescs=");
861         FDSpecifierStr += utostr(ToExecutor[ReadEnd]);
862         FDSpecifierStr += ',';
863         FDSpecifierStr += utostr(FromExecutor[WriteEnd]);
864         FDSpecifier = std::make_unique<char[]>(FDSpecifierStr.size() + 1);
865         strcpy(FDSpecifier.get(), FDSpecifierStr.c_str());
866     }
867
868     char *const Args[] = {ExecutorPath.get(), FDSpecifier.get(), nullptr};
869     int RC = execvp(ExecutorPath.get(), Args);
```

```
61 int openListener(std::string Host, std::string PortStr) {
62     addrinfo Hints{};
63     Hints.ai_family = AF_INET;
64     Hints.ai_socktype = SOCK_STREAM;
65     Hints.ai_flags = AI_PASSIVE;
66
67     addrinfo *AI;
68     if (int EC = getaddrinfo(nullptr, PortStr.c_str(), &Hints, &AI)) {
69         errs() << "Error setting up bind address: " << gai_strerror(EC) << "\n";
70         exit(1);
71     }
72
73     // Create a socket from first addrinfo structure returned by getaddrinfo.
74     int SockFD;
75     if ((SockFD = socket(AI->ai_family, AI->ai_socktype, AI->ai_protocol)) < 0) {
76         errs() << "Error creating socket: " << std::strerror(errno) << "\n";
77         exit(1);
78     }
79
80     // Avoid "Address already in use" errors.
81     const int Yes = 1;
82     if (setsockopt(SockFD, SOL_SOCKET, SO_REUSEADDR, &Yes, sizeof(int)) == -1) {
83         errs() << "Error calling setsockopt: " << std::strerror(errno) << "\n";
84         exit(1);
85     }
86
87     // Bind the socket to the desired port.
88     if (bind(SockFD, AI->ai_addr, AI->ai_addrlen) < 0) {
89         errs() << "Error on binding: " << std::strerror(errno) << "\n";
90         exit(1);
91     }
92
93     // Listen for incoming connections.
94     static constexpr int ConnectionQueueLen = 1;
95     listen(SockFD, ConnectionQueueLen);
```

The plan

- A MemoryMapper interface with implementations based on
 - Shared memory
 - When both executor and controller process share same physical memory
 - Regular memory allocation APIs
 - When the resultant code is executed in the same process
 - Useful for unit tests
 - EPC
 - Required when the executor and controller process run with different physical memory
 - Resultant code is transferred to the executor process over the EPC channel
- A JITLinkMemoryManager implementation that can use any MemoryMapper
 - It will allocate large chunks of memory using MemoryMapper and divide into smaller chunks
 - Better support for small code model by keeping everything close in memory

Current Progress

- MemoryMapper interface in review
- InProcessMemoryMapper implementation using sys::Memory APIs in review
- SharedMemoryMapper needs to be adapted to new MemoryMapper interface design (Currently working)

Thank you