Robustness of Message Passing Systems

- System Upgrade and Evolution
- Error Handling
- Erlang

CC BY ND

System Upgrade and Evolution

- Message passing allows for easy system upgrade
  - Rather than pass messages directly to server, pass them via proxy
  - Proxy can load a new version of the server and redirect messages, without disrupting existing clients
  - Eventually, all clients are talking to the new server; old server is garbage collected
- Allows for gradual transparent system upgrade without disrupting service

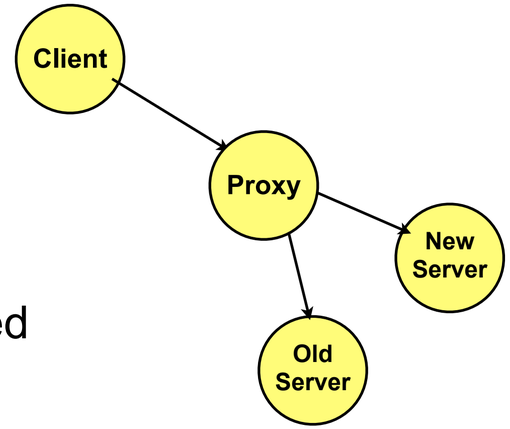

- Use of dynamic typing can make the upgrade easier
  - New components of the system can generate additional messages, which are ignored by old components
  - Supervisor hierarchy allows system to notice if components fail, and fallback to known good version
  - Backwards compatible extensions are simple to add in this manner

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Error Handling

- The system is massively concurrent – errors in one part can be handled elsewhere
- Error handling philosophy in Erlang:
  - Let some other process do the error recovery
  - If you can't do what you want to do, die
  - Let it crash
  - Do not program defensively

J. Armstrong, "Making reliable distributed systems in the presence of software errors", PhD thesis, KTH, Stockholm, December 2003, http://www.sics.se/~joe/thesis/armstrong_thesis_2003.pdf

- Be concerned with the overall system reliability, not the reliability of any one component

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Let It Crash

- In a single-process system, that process must be responsible for handling errors
  - If the single process fails, then the entire application has failed
- In a multi-process system, each individual process is less precious – it's just one of many
  - Changes the philosophy of error handling
  - A process which encounters a problem should not try to handle that problem – instead, fail loudly, cleanly, and quickly "let it crash"
  - Let another process cleanup and deal with the problem
  - Processes become much simpler, since they're not cluttered with error handling code

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Remote Error Handling

- How to handle errors in a concurrent distributed system?
  - Isolate the problem, let an unaffected process be responsible for recovery
  - Don’t trust the faulty component
  - Analogy to hardware fault tolerance

- Processes are linked, runtime is set to trap errors and send a message to the linked process on failure
  - e.g., process PID2 has requested notification of failure of PID1; runtime sends an “EXIT” message on failure, to tell PID2 that PID1 failed, and why
  - Process PID2 then restarts PID1, and any other dependent processes

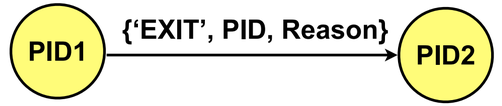

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Supervision Hierarchies

- Organise problems into tree-structured groups of processes, letting the higher nodes in the tree monitor and correct errors in the lower nodes
  - Supervision trees are trees of supervisors – processes that monitor other processes in the system
  - Supervisors monitor workers – which perform tasks – or other supervisors
  - Workers are instances of behaviours – processes whose operation is characterised by callback functions (i.e., the Erlang equivalent of objects)
    - E.g., server, event handler, finite state machine, supervisor, application

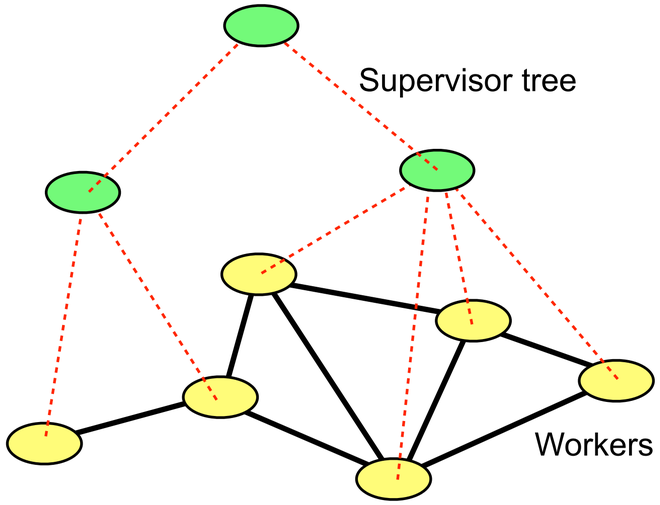

- Abstract common behaviours into objects
- Workers managed by supervisor processes that restart them in the case of failure, or otherwise handle errors

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Robustness of Erlang Systems

- Example: Ericsson AXD301 ATM switch
  - Dimensioned to handle ~50,000 simultaneous flows with ~120 in setup or teardown phase at any one time; processes traffic at 160Gbps (16 x 10Gbps links); written in Erlang
  - 99.9999999% reliable in real-world deployments on 11 routers at major Ericsson customer (0.5 seconds downtime per year)
    - Failures do occur, but are recovered by supervision hierarchy and distributed error recovery
    - Employs restart-and-recover semantics per-connection so failures disrupt only one connection out of thousands



Images from: S. Blau, J. Rooth, J. Axell, F. Hellstrand, M. Buhrgard, T. Westin, and G. Wicklund, “AXD 301: A new generation ATM switching system”, Ericsson Review, 1998.

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Further Reading

- J. Armstrong, “Erlang”, Communications of the ACM, 53(9), September 2010, DOI:10.1145/1810891.1810910

- **Does the programming model make sense?**
- **Does the reliability model (“let it crash”) make sense?**



contributed articles

DOI:10.1145/1810891.1810910

The same component isolation that made it effective for large distributed telecom systems makes it effective for multicore CPUs and networked applications.

BY JOE ARMSTRONG

Erlang

ERLANG IS A concurrent programming language designed for programming fault-tolerant distributed systems at Ericsson and has been (since 2000) freely available subject to an open-source license. More recently, we've seen renewed interest in Erlang, as the Erlang way of programming maps naturally to multicore computers. In it the notion of a process is fundamental, with processes created and managed by the Erlang runtime system, not by the underlying operating system. The individual processes, which are programmed in a simple dynamically typed functional programming language, do not share memory and exchange data through message passing, simplifying the programming of multicore computers.

Erlang[2] is used for programming fault-tolerant, distributed, real-time applications. What differentiates it from most other languages is that it's a concurrent programming language; concurrency belongs to the language, not to the operating system. Its programs are collections of parallel processes cooperating to solve a particular problem that can be created quickly and have only limited memory overhead; programmers can create large numbers of Erlang processes yet ignore any preconceived ideas they might have about limiting the number of processes in their solutions.

All Erlang processes are isolated from one another and in principle are "thread safe." When Erlang applications are deployed on multicore computers, the individual Erlang processes are spread over the cores, and programmers do not have to worry about the details. The isolated processes share no data, and polymorphic messages can be sent between processes. In supporting strong isolation between processes and polymorphism, Erlang could be viewed as extremely object-oriented though without the usual mechanisms associated with traditional OO languages.

Erlang has no mutexes, and processes cannot share memory.[a] Even within a process, data is immutable. The sequential Erlang subset that executes within an individual process is a dynamically typed functional programming language with immutable state.[b] Moreover, instead of classes, methods, and inheritance, Erlang has modules that contain functions, as well as higher-order functions. It also includes processes, sophisticated error handling, code-replacement mechanisms, and a large set of libraries.

Here, I outline the key design criteria behind the language, showing how they are reflected in the language itself, as well as in programming language technology used since 1985.

Shared Nothing

The Erlang story began in mid-1985 when I was a new employee at the Ericsson Computer Science Lab in Stock-

a The shared memory is hidden from the programmer. Practically all application programmers never use primitives that manipulate shared memory; the primitives are intended for writing special system processes and not normally exposed to the programmer.

b This is not strictly true; processes can mutate local data, though such mutation is discouraged and rarely necessary.

68 COMMUNICATIONS OF THE ACM | SEPTEMBER 2010 | VOL. 53 | NO. 9

CC BY ND

Summary

- Concurrency and memory models
- Transactions
- Message Passing

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow