Complexity and the Evolution of Computing: Metaphors for Managing Emergent Systems

Steve Burbeck, IBM

September 28, 2004

TTI/Vanguard Meeting, Los Angeles

- Carver Mead said years ago that “engineers would be foolish to ignore the lessons of a billion years of evolution.”
- He was talking about designing VSLI chips. His observation is even more applicable today as our computing systems seem to have a life of their own.
- But biological analogs are plentiful. Which ones will illuminate our struggle with complex systems?

I Suggest Multicellular Computing

- Assume computing is at a stage analogous to the transition between single cell and multicellular life
  - Computing is not so much about what goes on inside the CPU, but what goes on between them
- My email tag line has long made this assertion. Only a handful of people have ever commented on it
  - Perhaps it's obviously silly or irrelevant
  - Perhaps it's obviously true, but so what?
  - Perhaps few people know much about that transition
- The assumption is *not* amenable to proof or disproof
  - We don't have the perspective to see how computing evolves today
  - We don't really know how multicellular life evolved far in the past
  - No two complex systems evolve exactly the same. Why assume any evolutionary process gives insights for another?
- But...we will plunge on regardless

Starting First with Conclusions

- Multicellular life relies on (at least) four organizing principles:
  - specialization - about 250 different types of cells in humans
  - Cells communicate with each other via messenger molecules, *never* DNA. The “meaning” of cell-to-cell messages is determined by the receiving cell
  - Groups of cells build extracellular structures which, in turn, define the persistent “self.” That is, “selfness” resides as much in the extracellular matrix as in the cells
  - Every *healthy* cell is prepared to commit suicide
- All have implications for organizing multicellular computing systems

Complexity Begets Emergence

A flock of starlings

- Whenever sufficiently large numbers of elements interact dynamically, autocatalytic sets (sets of elements with self reinforcing dynamics) emerge.
- Each new level is qualitatively different from the one that gave rise to it
- Cause and effect crosses levels in ways that cannot be predicted, and can be analyzed only with difficulty

•Complexity grows (more elements, more interactions)
•Autocatalytic sets form
•These sets become the elements of the next level
•Repeat at next level
quarks...atoms...molecules...organelles...cells...
tissues...organs...organisms...ecologies...

The Internet Begets Multicellular Computing

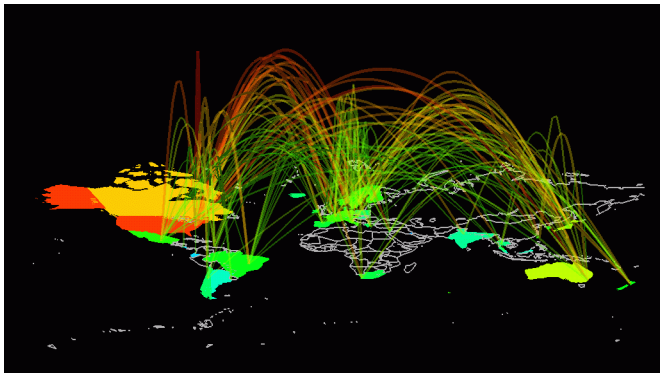

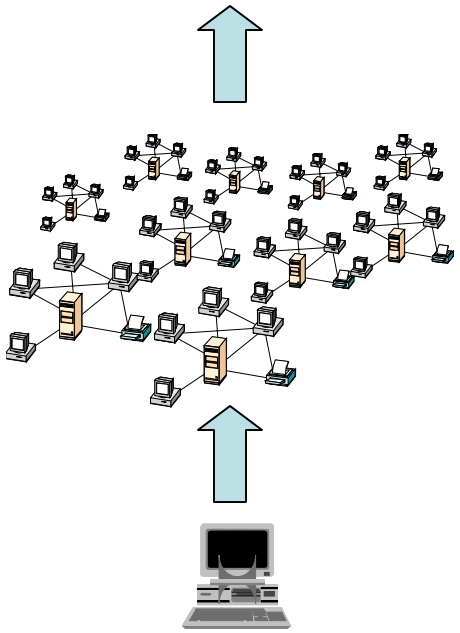

- Lots of machines connected to one another via the Internet
- Lots of different kinds of machine to machine interaction
- So, sets of machines form to collaborate for various purposes
  - Corporate infrastructures
  - P2P filesharing and VOIP networks
  - Grids

The Internet Also Begets Co-evolving Predators

- Viruses, worms, spam, spin, spit, popups, popunders, phishing, ...
- A large monoculture such as Windows/IE attracts the most predators
  - Nearly 5,000 new Windows viruses and worms were documented in the first half of the 2004, up from 1000 the year before

Multicellular Organisms as Metaphor

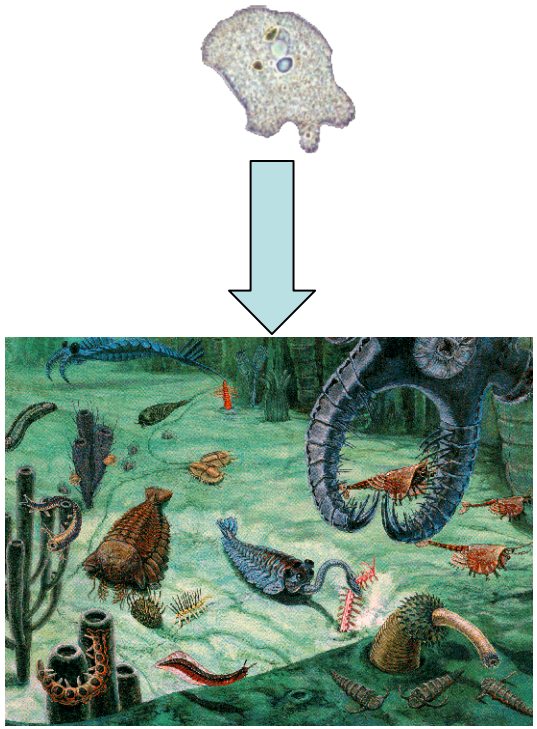

- The evolution of multicellular life faced ***four problems*** similar to those we face in multicellular computing
  - overly complex individual cells make uncomfortable bedfellows
  - need “safe” communication and collaboration
  - how to distinguish self from non-self
  - how to resist microbial and viral attack

Biofilms: Training Wheels for Multicellular Life

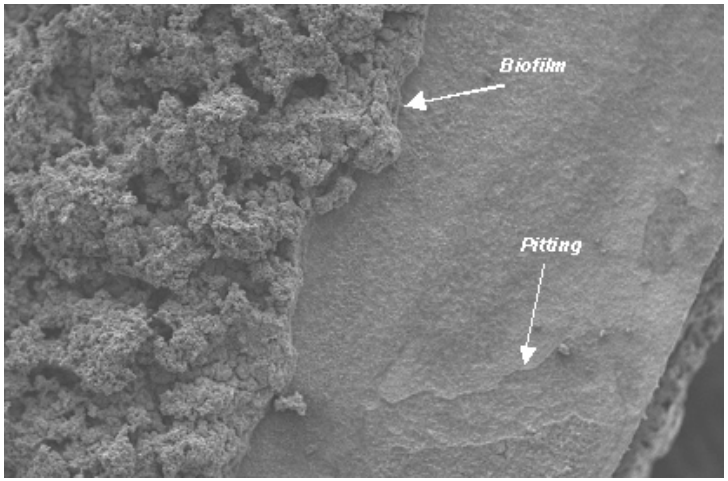

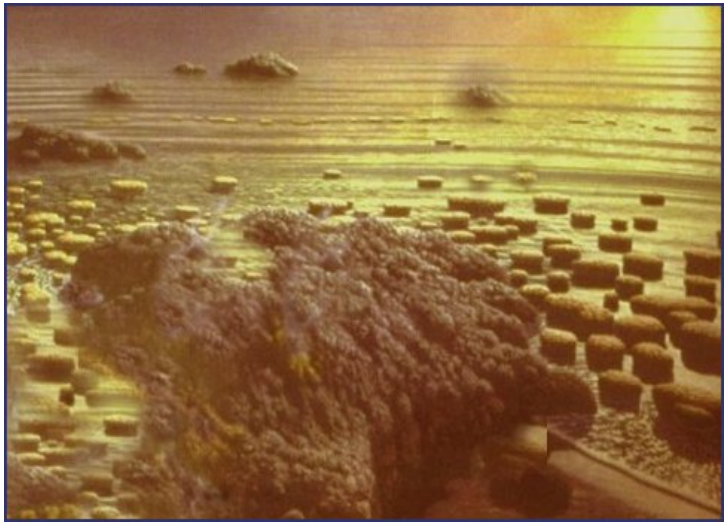

- Biofilms are surface-attached accumulations of collaborating cells encased in extracellular polymeric substances secreted by the cells
  - for example, Stromatolites are the fossilized remains of Precambrian colonies of cyanobacteria
- Biofilm cells specialize in temporary, context specific, ways
- They communicate by protein messages
- The colonies are fundamentally stigmergic
- Apoptosis appears to have evolved coincident with early bacterial colony organisms

Issue 1: Specialization and Differentiation

- Single cell organisms evolve to be increasingly complex
  - Each cell is on its own so the more adaptable and multi-function, the better. Species like paramecia and amoeba are the “top predators”
  - The burden of this complexity is slower growth due to increased metabolic needs
- Multicell organisms specialize behavior of individual cells
  - Cells in biofilms and other cooperative bacterial colonies retain full functional capability. Their specialization is temporary.
  - Cells in “true” multicellular organisms permanently differentiate (with the exception of stem cells). Reversal of differentiation, neoplastic transformation, is an early indication of cancer.

Differentiation in Multicellular Computing

- Today's multicellular computing systems are much more akin to biofilms than to fully differentiated multicellular organisms.
  - Few of today's individual computers are specialized, even temporarily, for exposure to the Internet
  - Loaded with general-purpose functions (e.g., display of jpg images) that can become an entry point for viral or worm attack
- Multicellular life suggests that using the same general-purpose machine for all functions will eventually fade away
  - We should look to a similar approach

Issue 2: Safer Communication and Collaboration

- Complex collaboration between cells requires information-rich messages
- Two sorts of complex molecules offer themselves. Both are used in certain circumstances
  - DNA offers highest information content. Used when information must be passed on by the receiver to its offspring
  - Proteins perhaps offer less information content but cannot be “executed” by the cell replication machinery...much safer

Bacteria Exchange DNA

- DNA exchange “reprograms” the receiving cell
- The receiver has no way to decide to accept the new behavior or not
- Comparable to transmitting a .exe file ... Or a virus

The Multicellular Taboo

- Multicellular organisms ***do not exchange DNA*** outside of sexual reproduction
- The rule against genetic transfer is so universally obeyed in multi-cellular organisms that Loewenstein calls it "… the taboo of intercellular transfer of genetic information"
  - See The Touchstone of Life, Werner Loewenstein, Oxford University Press, New York, 1999 [p. 277]
- Multicellular organisms communicate by polymorphic messenger molecules
  - primarily proteins because of their high information content and easy control via DNA binding proteins

Prevent Promiscuous Computers

- Many PCs allow easy download of executable code
- This unsafe sex invites viral and worm infections
- Active-X -- the Worst Offender
  - “...hackers continue to find and exploit security holes in Explorer. Many of them take advantage of Explorer's ActiveX system, which lets Web sites download and install software onto visitors' computers, sometimes without users' knowledge.” (Slate, June 30)
  - Other browsers, e.g., Mozilla’s Firefox, protect against ActiveX and other well-known infection paths.
- We must adopt the multicellular life taboo
  - messages should be XML (or the like), never executable code
  - Depends upon presence of interpreters (data becomes code with the right interpreter)

Issue 3: Emergence of “Self” in Multicellular Organisms

- We tend to think of the “self” as having to do with the immune “identity” of individual cells
  - but the immune systems did not evolve until after multicellular life (perhaps with the early cniderians such as the hydra)
  - And, the immune system is not sufficient
    - blood transfusion -- when immunilogically identified cells leave the physical body, they no longer are “self”
    - kidney transplant - when an “other” organ joins the body, it *becomes* self

Stigmergy Structures

- Stigmergy refers to activity organized over time by creation of and interaction with an external structure -- like a termite mound, or a honeycomb
- A biofilm exudes and is organized by its extracellular gel
- A higher organism is organized by its skeleton, muscle, sinew, all of which are ever-changing structures created by and intertwined with specialized cells
- “Self” is properly associated with the external stigmergy structure more than with the individual cells

Stigmergy in Computing

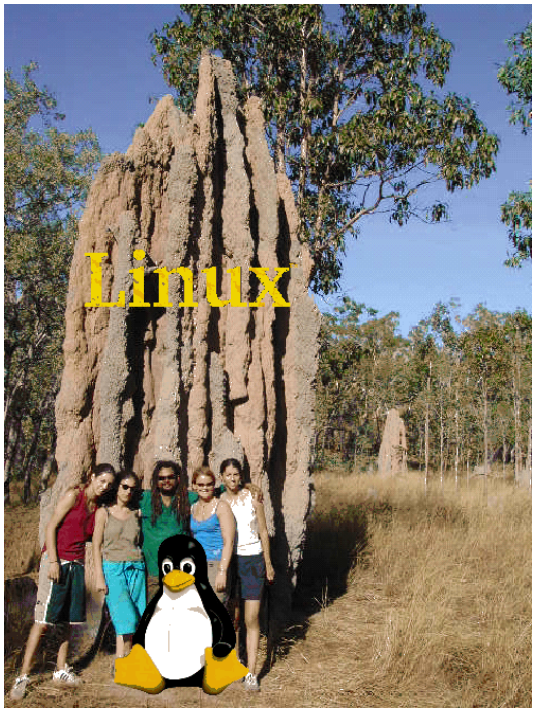

A software stigmergy structure

- Corporate networks, databases, email systems
- Linux, where programmers build an OS in the stigmergy structure of the CVS code tree repository
- P2P file-sharing or VOIP systems where the “cells” cooperatively construct the “presence registry” in real-time and use it to organize search and transfer

Issue 4: Defending the Self

- Programmed Cell Death (PCD), also called Apoptosis turns out to be everpresent and powerful
- It is defense by suicide

The Emergence of Apoptosis

- Apoptosis is a prerequisite for multicellular life
  - Apoptosis evolved coincident with the first types of multicellular life - colonies of bacteria - although it has become more elaborate in higher-level organisms
  - It evolved to deal with the sorts of issues that plague multicellular organisms but not single cell organisms -DNA replication errors, viral infection, and escape from cell differentiation (neoplastic changes).
  - And it solves those issues from a multicellular perspective - sacrificing the individual cell for the good of the multicellular organism.

Apoptosis (Programmed Cell Death)

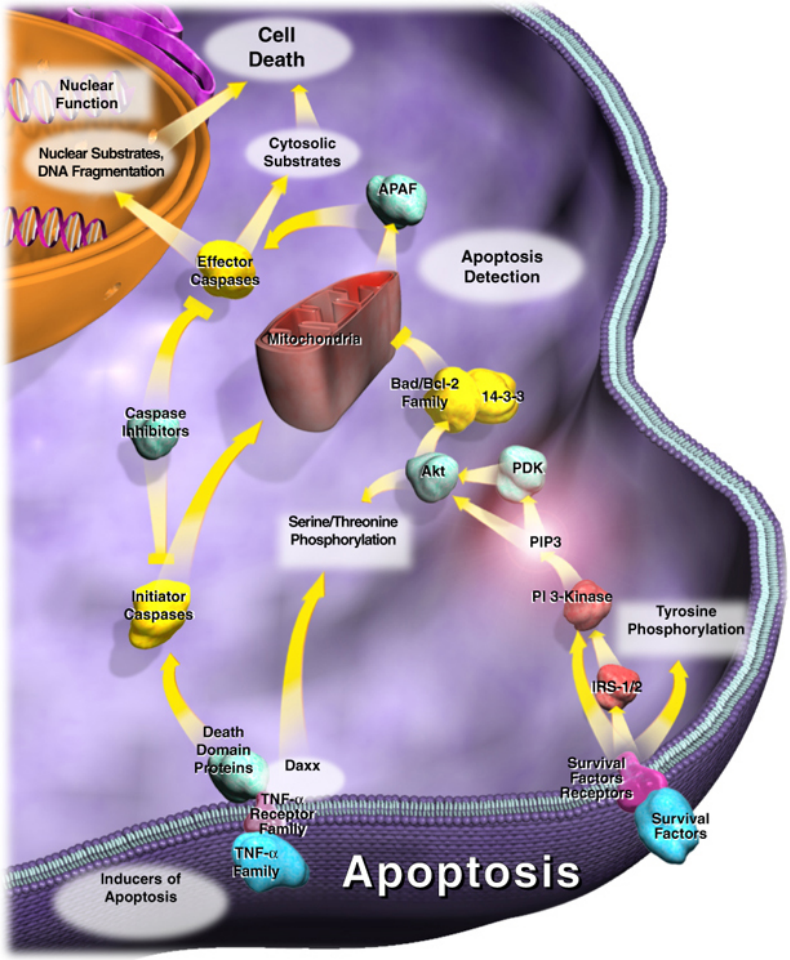

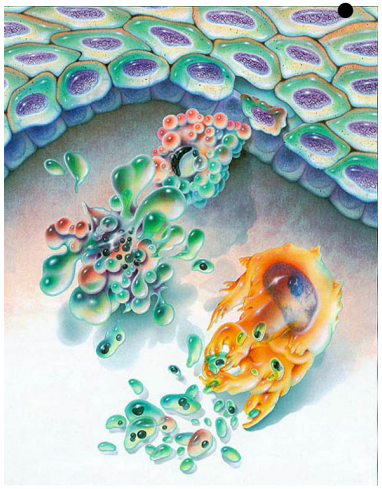

- A carefully choreographed process
  - The cells shrink
  - Develop bubble-like blobs on their surface
  - Degrade the chromatin (DNA and protein) in their nucleus
  - Break down mitochondria
  - Break the cell into small, membrane-wrapped, fragments
  - Expose the phospholipid phosphatidylserine, normally hidden within the outer membrane
  - This is bound by receptors on phagocytic cells like macrophages and dendritic cells which then engulf the cell fragments
  - The phagocytic cells secrete cytokines that inhibit inflammation

Triggers for Apoptosis

- The balance between:
  - the receipt of positive signals; that is, signals needed for continued survival and, for many, continued adhesion to the surface on which they are growing.
    - growth factors such as EGF (Epidermal Growth Factor)
    - Interleukin-2 (IL-2), an essential factor for the mitosis of lymphocytes
  - the receipt of negative signals.
    - increased levels of oxidants within the cell
    - damage to DNA by these oxidants or other agents like
      - ultraviolet light and other ionizing radiation
      - chemotherapeutic drugs or other mutagenic chemicals
    - accumulation of proteins that failed to fold properly into their proper tertiary structure
    - molecules that bind to specific receptors on the cell surface and signal the cell to begin the apoptosis program. These death activators include:
      - Tumor necrosis factor (TNF-alpha) that binds to the TNF receptor;
      - Lymphotoxin (also known as TNF-beta) that also binds to the TNF receptor;
      - Fas ligand (FasL), a molecule that binds to a cell-surface receptor named Fas (also called CD95)

The “Blue Screen of Death” as Apoptosis

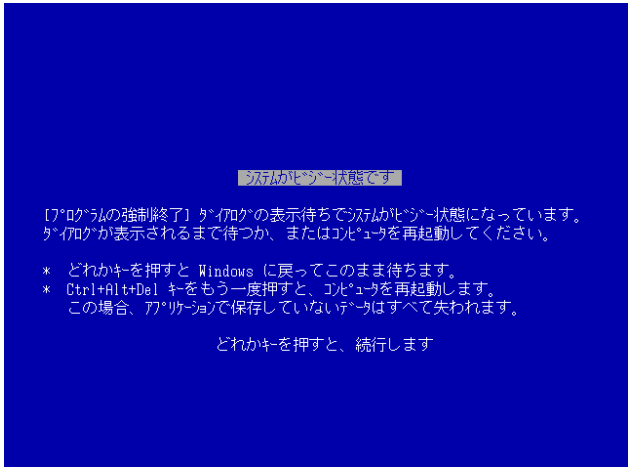

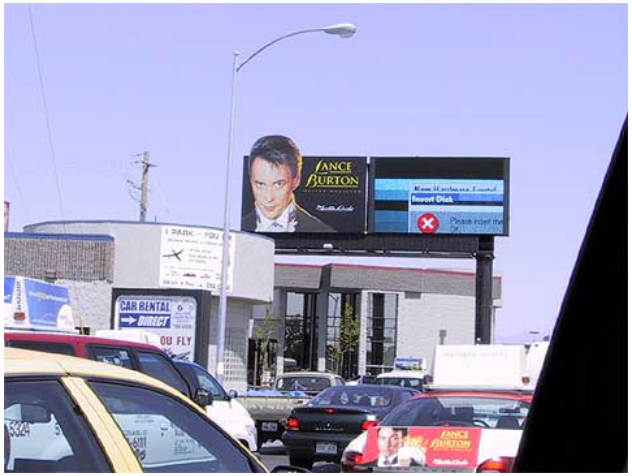

- Apoptosis is preferable to necrosis
- While no one welcomes the BSoD, we tend to forget that it is a “controlled” death as opposed to a complete lock-up
- It arises when the operating system realizes that something serious has been corrupted

Summary of Derived Principles

- Seek specialization and avoid monoculture
  - move complexity from within individual cells to the diversity of the collaborating group of cells
  - an infection, like a cold, should be able to infect only a small proportion of the cells
  - use autonomic computing techniques to manage that diversity
- Base collaboration on messaging, never code transfer
  - Be stingy with language interpreters and be careful what powers you allow them
- “Self” is a property of attachment to the stigmergic structure, not of the individual cell
  - Keep your eye on the persistent, perhaps boring, data and networks, not the flashier computing stuff
- One cell is never more important than the multicellular “self”
  - Apoptosis is your friend. If the cell strays, kill it. If the hand offends, cut it off

Barriers to Acceptance

- The industry presumes lots of cheap powerful general purpose machines
  - The economics favor that model
- More specialization may pose more problems for configuring and managing the systems (hence more need for autonomic computing)
- Finding the right mechanisms and triggers for computer apoptosis will take time