# 11 How I Invented a New Internet

Lecture 180120-1
https://www.youtube.com/watch?v=47RzSvweTOE


Philip Emeagwali Lecture 180120-1
Visit http://emeagwali.com for complete transcripts of 100+ lectures.
Video:
https://www.youtube.com/watch?v=47RzSvweTOE

https://YouTube.com/emeagwali

Podcast: https://SoundCloud.com/emeagwali

# 11.1 How I Named My New Internet

## 11.1.1 Naming of My New Internet

Scientific knowledge
is the first son of God.
Science **pre-existed** before humanity
and before our planet, the Earth,
was formed 4.6 billion years ago.
Back in 1989,
one of the science news headlines
was that an African Supercomputer Wizard
in the United States
had experimentally discovered
how and why parallel processing
makes modern computers **faster**
and makes the new supercomputer
the **fastest**
and invented
how and why to use
that **new supercomputer knowledge**

to build a **new supercomputer**
that encircled the globe
in the way the internet does.
I am that African supercomputer scientist
who was in the news
back in 1989.
I was in the news
for experimentally discovering
that parallel processing
is an entirely new way of supercomputing
**across** thousands or millions or billions
of processors.
Parallel processing
is defined as the technique
of fastest supercomputing
that is fastest
by computing many things
**at once,** or in parallel,
instead of computing only one thing
**at a time,** or in sequence.
Prior to my 1989 invention,
parallel processing was widely caricatured
and rejected

as a huge waste of everybody's time.
Parallel processing
was rejected for four reasons.
The first reason
the parallel processing supercomputer
was rejected
was because supercomputing in parallel
had performance problems.
That is, in the 1980s and earlier,
parallel processing supercomputers
could not compute faster than
sequential processing supercomputers.
The second reason
the parallel processing supercomputer
was rejected
was because it was physically impossible
to invent
how to harness
64 binary thousand processors
and harness them
to compute together to solve
any of the twenty toughest problems
arising in supercomputing.

Those **extreme-scale** problems
were called the twenty Grand Challenges
of supercomputing.
The third reason
the parallel processing supercomputer
was rejected
was that programming supercomputers
to solve a system of coupled, nonlinear,
time-dependent, and state-of-the-art
partial differential equations
of a new calculus
made research computational
mathematicians
deeply uncomfortable.
In particular, to **parallel process**
via emails
sent to and from
sixteen-bit long email addresses
and to **parallel process**
the most dense, abstract,
and impenetrable equations
and to **parallel process**
their algebraic approximations

and to **parallel process**
their floating-point arithmetical calculations
that must be executed across
sixteen times
two-to-power sixteen,
or across one binary million, email wires
is like dancing in the fire.
The fourth reason
the parallel processing **supercomputer**
was rejected
was that I, its discoverer,
was black and African.
My research and experimental discovery
of parallel processing
was not taken seriously
in the late 1970s and early 1980s.
My 1,057 page research report
on the massively parallel processing
supercomputer
was rejected six times
and rejected by three universities
and rejected by scientific journals
before it was eventually accepted

by the supercomputer community.
In the 1980s, the massively
parallel processing supercomputer
was unfathomable
and for that reason
a president of an American university
that had an annual research expenditure
of one billion dollars
and his five supercomputer experts
threw my one thousand
and fifty-seven [1,057]-page
supercomputer research report
into the trash.
When a newspaper journalist
writing about my invention
came to interview
those five supercomputer experts
they couldn't do the interview.
The reason was that
they never read or understood
my supercomputer research report.
So I was not taken seriously
until The Computer  Society

of the IEEE—The Institute of Electrical
and Electronics Engineers—
gave me the top prize in supercomputing.
To put my **dilemma** in context,
back in the 1980s,
it was impossible
for an all-white scientific jury
to give me the top award
in computer science.
The award committees
asked for my photograph
or insisted on a face-to-face interview
that will reveal
the fact that I am black and African.
In the 1980s,
only one award committee
did not demand my photograph.
I won that award
and it made the news headlines
that a black African
had won the top prize
in supercomputing.
The controversy

prompted the award committee
to change their rules
and to demand a face-to-face lecture
that, in turn, made it impossible
for other black supercomputer scientists
to win the top prize in supercomputing.
To this day,
the color of my skin
gets more attention
than the solution of my equations.
In the 1980s and earlier
and in the United States,
white research mathematicians
did not attend research seminars
given by black research mathematicians.

## 11.1.2 Why I Won the Top Prize in Supercomputing

When the Computer Society
of The Institute of Electrical
and Electronics Engineers
gives its top award to a supercomputer

inventor,
computer scientists
and the 450 thousand members
of The Institute reads about it.
After my widely-rejected
experimental discovery
was accepted and validated
by The Computer Society,
the naysaying vector processing
supercomputer scientists
that—at that time—did not believe in
parallel processing supercomputers
saw The Computer Society's endorsement
of my experimental discovery
as a vote of confidence
on massively parallel processing
supercomputers.
The public saw the news headlines
on the African supercomputer wizard
that won top US prize
as a vote of confidence
on **Philip Emeagwali**.
In the decades of the 1960s

through ‘80s,
parallel processing
was the subject of a titanic battle
between the **majority**
who believed that all supercomputers
should be powered by
a single, isolated processor
and the **minority**
who believed that all supercomputers
should be powered by
an ensemble of thousands of processors.
That was the reason
only one computational mathematician
attended my public lecture
on **parallel processing**
that took place in November 1982
and took place in a lecture auditorium
that was a short walk
from The White House, Washington, D.C.
Nine years later, my lecture
on **parallel processing** supercomputing
that I gave on July 8, 1991
in Washington, D.C.

was before a **standing** room only audience
of research computational mathematicians
that were attending
the largest international congress
of mathematics.
That audience—that was similar
to the one of nine years earlier
that humiliated, ridiculed,
and rejected my experimental discovery
of parallel processing—gave me
a **standing ovation**.

### 11.1.3 The Free Performance Lunch is Over

After my experimental discovery
of how and why parallel processing
makes modern computers **faster**
and makes the new supercomputer
the **fastest**
my telephone began to ring off the hook.
It seemed like every other
research computational scientist

wanted to become my new best friend
and my new scientific collaborator.
So, I was not surprised
when **Steve Jobs**
tried to reach me by telephone
in about June 1990.
**Steve Jobs**
wanted to know how he could harness
the power of parallel processing
to process images
and to do so **faster**.
To put things in context,
back in June 1990, **Steve Jobs**
was depressed and devastated
because he was unceremonious removed
from Apple Corporation,
the company that he started.
Looking for a new direction,
**Steve Jobs**
was intrigued by my experimental discovery
of how and why parallel processing
**across** a global network of
65,536 processors,

or **across** a new internet,
reduced 65,536 days, or 180 years,
of **time-to-solution**
on only one processor
that is not a member
of an ensemble of processors
to just one day of **time-to-solution**
**across** a new internet
that is a global network of
65,536 commodity-off-the-shelf processors.
Fast forward eighteen years,
to June 9, 2008,
**Steve Jobs**
told the opening session
of Apple's Worldwide Developers Conference
in San Francisco, California
that parallel processing
is still very challenging.
As reported, the following day,
in the June 10, 2008 issue
of the *New York Times*,
**Steve Jobs**
told Apple's Worldwide Developers

that: [quote]
PROP ALERT
“The way the processor industry
is going
is to add more and more cores,
but nobody knows
how to program
those things,”
**Steve Jobs** said.
And he continued:
“**I mean, two, yeah;**
**four, not really;**
**eight, forget it.”**
[unquote]
I experimentally discovered
how and why
massively parallel processing
is at the heart of the fastest supercomputer.
I experimentally discovered
that massively parallel processing
is a necessary condition
for the fastest supercomputers.

Historically, we never had
new supercomputers without
experimentally discovering
faster supercomputer speeds.
To achieve grand wizardry
in fastest massively parallel supercomputing
requires the **visceral** understanding
that the massively parallel supercomputer
is not a new computer, *per se*.
I experimentally discovered
that my new and massively
parallel supercomputer
that I visualized
as a global network of
65,536 processors
is a small internet, *de facto*.

### 11.1.4 My Naming Convention Across a New Internet

I invented a new internet
that was defined and outlined
by a new global network of

65,536 tightly-coupled processors.
And I invented
how to use that new internet
to make modern computers **faster**
and to make the new supercomputer
the **fastest**
and how and why to use
that **new supercomputer knowledge**
to build a **new supercomputer**
that encircled the globe
in the way the internet does.
I will take a **retrospective** look
on my early years,
or the sixteen years
onward of June 20, 1974.
I will look back on how I named
each processor within my new internet
that is a global network of
64 binary thousand processors.
And how I invented
that new internet
to be a massively parallel supercomputer.

I experimentally invented
how to assign a unique string of
sixteen zeroes and ones
and assign each string
as the sixteen-bit name
of each of my two-to-power sixteen codes
that had a one-to-one correspondence
with my 64 binary thousand processors.
So after several years
of hands-on, direct programming
of an ensemble of 64 binary thousand
commodity-off-the-shelf processors
that I visualized
as my new internet,
I became known to programmers
in the supercomputing community
as the [quote unquote]
"go-to" person.
Research supercomputer scientists
that sought answers to questions
on message passing came to me.
Vector processing supercomputer
programmers—who at that time

presumed that I was in the
Los Alamos National Laboratory
in Los Alamos, New Mexico, United States
that was the supercomputing capital
of the world
emailed me
when they wanted to learn
how to program
the message-passing ensemble of processors
that everybody hated
and that everybody ridiculed, mocked,
and disrespected
as a huge waste of everybody's time.
That unique message-passing
supercomputer experience
that I gained in the 1980s
was the reason,
I was appointed in the early 1990s
as the Distinguished Lecturer
of the two leading computer societies
in the world,
namely, the Association
for Computing Machinery

and The Computer Society
of the Institute of Electrical
and Electronics Engineers.

## 11.1.5 Father of the Internet

A twelve-year-old writing a school report
on "**Philip Emeagwali**"
asked me:
"Why are you called
the father of the Internet?"
I answered:
"**The internet**
**has many fathers and mothers,**
**uncles and aunts.**
**But I am the only father**
**of the Internet**
**that invented a new internet.**
**I am the only father**
**of the modern supercomputer**

**who was profiled**
**in major U.S. newspapers**
**and who was credited**
**for the invention**
**of the massively parallel processing supercomputer**
Back in June 1974,
I conceived 64 thousand computers
around the Earth
that comprised of
a new global network of computers.
I conceived that new global network
as used to forecast the weather.
But it took me fifteen years,
onward of June 1974,
to invent how to harness
that new global network of computers
and harness it to forecast the weather."
In 1989, it made the news headlines
that an **African supercomputer wizard**
in the United States

had won the top prize
in the field of supercomputing.
That **African supercomputer wizard**
was in the news for inventing
how to harness a new internet
that is a new global network of
65,536 tightly-coupled
commodity-off-the-shelf processors.
Each processor
was akin to a tiny computer.
I am that African
internet scientist
who was in the news
back in 1989.
I won the top prize in supercomputing
because I experimentally discovered
how an **ensemble** of 65,536,
or two-raised-to-power-16,
commodity-of-the shelf processors
could be **assembled**
as the building blocks

of a new supercomputer
and harnessed to become
the world's fastest supercomputer,
*de facto*.
I was in the news
because I invented
how to synchronously communicate
and how to simultaneously compute
and how to communicate
and compute **together**
and how to do both
as one seamless, cohesive unit.
That cohesive unit
was my new supercomputer *de facto*.
That cohesive unit
was defined around
a sixteen-dimensional hyperball
that is a new internet, by definition.
That cohesive unit
was the supercomputing machinery
that I used to **send** and **receive** emails

to and from
65,536, or two-to-power sixteen
sixteen-bit long
email addresses.
Each of those 64 binary thousand
email addresses
was a unique string of
16 zeroes and ones.
Back in the 1980s,
I emailed 64 binary thousand
computer codes
to as many processors.
Each of those computer codes
solved initial-boundary value problems
with each problem's governing
partial differential equations
of modern calculus
and with each equation's specified
initial and boundary conditions.
What made the news headlines
was that I solved

24 million equations of algebra.
That was a world record
in algebra
back in 1989
as well as my contribution
to algebra.
Each equation of algebra
that I solved
was restating the Second Law of Motion
that was at the physics core
of the computational physics model
that I executed
within each processor.
I was in the news in 1989
for inventing
how to solve problems
in extreme-scale algebra
and for inventing how to solve them
**across**
a new internet
that is a new global network of

65,536 tightly-coupled processors
with each processor
operating its own operating system
and with each processor
having its own dedicated memory
that shared nothing with each other.
I was in the news in 1989
for inventing that new internet
and for inventing it
as a new supercomputer, *de facto*.
What made the news headlines
was that I synchronously communicated
via emails to and from
**across** 65,536 cooperating processors
and that I simultaneously computed
at the speed of 47,303
calculations per processor
to compute
at the then unheard of
total speed of
3.1 billion calculations per second.

That experimental discovery
of massively parallel processing
changed the way we think about
the new supercomputer
that is the fastest computer
that will become
the computer of tomorrow,
if history repeats itself.
That experimental discovery
of massively parallel processing
garnered international headlines
and I the story teller
became the story
and the subject of school reports titled:
"The Contributions of **Philip Emeagwali**
to the Development of the Computer."
I was the first to experimentally discover
how and why parallel processing
makes modern computers **faster**
and makes the new supercomputer
the **fastest**

and how to use
that **new supercomputer knowledge**
to build a **new supercomputer**.
I experimentally discovered
massively parallel processing
and I did so
by solving a grand challenge problem
that the United States government
defined as one of its
twenty gold-ring problems
in supercomputing.
My experimental discovery
changed the way we looked at
the supercomputer.
Back in the 1970s and ‘80s,
I looked at the precursor
to the modern supercomputer
that I programmed
as a parallel processing **internet**
that was outlined and defined
by my 65,536

commodity-off-the-shelf processors.
Fast forward four decades,
the modern supercomputer
is a union of vast numbers of processors
that **communicate**
as a tightly-coupled Internet
that is outlined and defined
by millions upon millions of processors.
That experimental discovery
of the parallel processing internet
is the reason I am profiled in books on
the history of the Internet.

## 11.1.6 How I Named My New Internet

**The internet**
**has many fathers and mothers,**
**uncles and aunts.**
**But I am the only father**

**of the Internet**
**that invented a new internet.**
**I am the only father**
**of the computer**
**who was profiled**
**in major U.S. newspapers**
**and who was credited**
**for the invention**
**of the massively parallel processing**
**supercomputer**
Each of my 65,536
tightly-coupled processors
with each processor
operating its own operating system
and with each processor
having its own dedicated memory
that shared nothing with each other
encircled my new internet
and had its unique name.
That name was a unique string of
sixteen zeroes and ones.
I used a unique **binary reflected**

naming scheme
for each processor
that was within my global network of
65,536 processors.
My global network of processors
is a small internet.
That small copy of the internet
is one of the keys
to my experimental discovery
of how and why parallel processing
makes modern computers **faster**
and makes the new supercomputer
the **fastest**
and my invention
of how to use
that **new supercomputer knowledge**
to build a **new supercomputer.**
The experimental discovery
of massively parallel processing **across**
a new internet
is my contribution
to the development of **faster** computers
and the **fastest** supercomputer.

I experimentally discovered
how and why
the millions of processors
of a massively parallel supercomputer
can be harnessed to cooperatively
compute together
and to compute as one seamless,
cohesive unit
and to compute faster than
any serial or any vector
processing supercomputer.
I experimentally discovered massively
parallel processing
and I invented the technology
through my proper naming
of the processors within my internet.
I visualized my internet
as encircling a globe, or a hyperglobe,
in hyperspace.
That experimental discovery
is my contribution
to the development
of the first internet

that's the fastest supercomputer
and that massively parallel processed
**across** an ensemble of
65,536 cooperating processors.

## 11.1.7 The Modern Supercomputer

In the 1960s, '70s, and '80s,
parallel processing
was dismissed as a huge waste of
everybody's time.
In the most quoted scientific paper
in supercomputing
that was published in April 1967,
**Gene Amdahl**—the supercomputer scientist
of **Amdahl's Law** fame—wrote that
the maximum speed increase
that could be achieved
from harnessing
an ensemble of eight processors