A STATISTICAL MECHANICS OF DISTRIBUTED RESOURCE SHARING MECHANISMS *

BY
YECHIAM YEMINI
COMPUTER SCIENCE DEPARTMENT
COLUMBIA UNIVERSITY, NY, NY 10027

July 15, 1982

---

* This research was supported in part by the Defense Advances Research Projects Agency, Project N00039-82-C-0427

paper will appear in Proceedings of Infocom '83

A STATISTICAL MECHANICS OF DISTRIBUTED RESOURCE SHARING MECHANISMS*

Yechiam Yemini
Computer Science Department
Columbia University,
NY, NY 10027

Abstract

The problem of analyzing the performance of large-scale interacting distributed resource sharing mechanisms arises in packet broadcast networks as well as multiprocessor switching mechanisms, VLSI chip communications and distributed databases. Queueing theory has major limitations in analyzing such systems: it cannot handle tight interactions and dependencies and it requires a fine-grained analysis of the detailed dynamics of a system to obtain coarse-grained equilibrium results. This paper proposes an alternative approach to the problem based upon statistical mechanics. Using this approach it is possible to analyze the performance of intricate distributed resource sharing mechanisms with a relative ease. Moreover, the analogy to physical phenomena offers new performance measures (e.g., communication "Energy" and "pressure") and physical insights to the behavior of such systems.

1. Introduction

Consider a set of distributed agents sharing a distributed resource. Access to the resource is subject to mutual interference between the agents. Given some statistics of the generation and duration of service requests, it is required to compute such quantities as the throughput, delay and blocking of the system. Examples of such systems, to name a few, arise in packet radio networks where transmissions sharing the broadcast medium interfere with each other, in multiprocessor switches where transmissions through the switch block each other, in database systems where queries block each other by locking subsets of files and in VLSI on-chip communications.

Classical queueing theory suffers three major difficulties when it comes to analyzing large-scale interfering distributed resource sharing mechanisms. First, queuing theory is a fine-grained microscopic theory. It requires a detailed study of the evolution of each component in the system even when the only information of interest is some global steady-state averages. Second, queueing theory offers very little when it comes to analysis of tight interaction among queueing systems. The theory only works when somehow a large scale problem may be decomposed into a collection of independent single-queue problems. Modelling large-scale interacting systems in a way that the solution is decomposable has been a source for many interesting research results. However, when it comes to interference systems as in the above examples, decomposition normally fails to provide adequate solutions. Third, Queueing theory requires that one analyze the (Markov process) dynamics that lead to a steady-state in order to obtain steady-state solutions. This is similar to trying to solve the equilibrium behavior of a gas by tracking the dynamics of the molecular motions that lead to equilibrium. The physics of non-equilibrium and approach to equilibrium behavior is an order of magnitude more complex than a direct equilibrium solution. Similarly, one would like to have equilibrium analysis of large-scale distributed resource sharing mechanisms that does not require study of dynamics.

A natural alternative to queueing theory is statistical mechanics. Statistical mechanics provides the tools for coarse-grained analysis and for handling interactions and coupling between components. Statistical mechanics, in contrast to Markovian analysis, does not track the detailed motion of a system towards its steady-state; rather, it uses an interaction potential to derive the steady-state behavior directly.

The objective of this paper is to demonstrate applications of ideas and tools borrowed from statistical mechanics to the analysis of large-scale computer communication mechanisms. The models presented are simplified and only very simple tools of statistical mechanics are employed. This choice does not represent any limits of the approach, only our interest in providing a simple to read seminal paper which does not embark upon unduly complex analysis.

The idea of applying statistical-mechanics tools to the analysis of large scale communication systems was, to the best of the author's knowledge, first proposed by Benes [BENE 63]. Loosely speaking, the major result of that work was to demonstrate the use of maximum-entropy approach to complex queueing systems in circuit-switched networks. This work has been followed-up by others [FERD 70, SHOR 78, BARD 80] where it was applied to a number of classical queueing systems.

This paper focuses on interference models of distributed resource-sharing mechanisms. The major achievement of the paper is in developing a complete physical approach to the analysis of such systems, based on the analogy to lattice interaction models [RUEL 69]. We show how analysis can be liberated from the detailed dynamics and proceed, similarly [illegible] the paradigm of the working physicist, directly [illegible] equilibrium interference potential. We apply the [illegible] a few interfering resource-sharing mechanism[illegible] thermodynamical (i.e., macroscopic) [illegible]

*This research was supported in part by the Defense Advanced Research Projects Agency, project N00039-82-C-0427.

networks, analogous to energy, volume, pressure and other thermodynamical functions. We study the order-disorder behavior of networks as they are "cooled" and show how some networks "crystallize" when they are sufficiently "cold". We demonstrate critical phase-transitions arising in some packet broadcast networks. Among other results we show how classically difficult problems such as analyzing blocking properties of networks are greatly simplified, and how global "physical intuition" of networks may be developed.

2. A Model Of Distributed Interacting Resource-Sharing Mechanisms

Consider a resource shared by a set of distributed agents. At any moment of time a given agent is either idle or busy utilizing the resource. Assume further that agents may interfere with each other and that two interfering agents may not acquire simultaneous access to the resource. We call such a system an interference system. An interference system may be represented by an interference graph whose nodes represent the agents and whose edges represent mutually exclusive interference between the respective agents. We call such a graph the interference graph of the system. In a packet radio network or a multiprocessor switch the contending agents are transmissions accessing the shared communication medium. In a database system queries are the agents contending over the use of the shared database.

In view of our primary interest in computer communication mechanisms we shall use the terms "transmissions" and "communication medium" to refer to agents and resource respectively. Consider an interference system described by an interference graph $G \triangleq \langle N,E \rangle$; where N and E indicate the nodes and edges of G. To describe the statistical model of the generation and duration of transmissions, we assume that a node in the graph is marked with either 0 (for idle) or 1 (for active). We assume that:

- An idle node becomes active according to an exponential interarrival law with rate λ.
- A transmission lasts an exponentially distributed time with rate μ.
- A transmission is blocked if a neighboring nodes in the interference graph is active.

This simple model and the respective results may be extended and applied to more complex models of interaction and traffic generation. Such extensions will be discussed in a later section. For the time being, however, we confine ourselves to interference models described above.

With the above assumptions the evolution of the network is that of a spatial birth-death process over the interference graph. Let $\pi(A)$ represent the equilibrium probability of the set $A \subseteq V$ being active while $N \setminus A$ being idle, then $\pi(A)$ satisfies the following equilibrium equations:

$$\pi(A)[\lambda|\bar{A}^c| + \mu|A|] = \sum_{x \in \bar{A}^c} \mu\, \pi(A \cup \{x\}) + \sum_{x \in A} \lambda\, \pi(A \setminus \{x\}) \qquad (1)$$

Here $\bar{A}$ (read, the closure of A) is the set of vertices in A and those neighboring to vertices in A.

To solve equation (1), let us define a set of nodes in the interference graph as independent if no pair of nodes in the set are neighbors. Independent sets of nodes represent possible concurrent transmissions configurations. Let $\mathfrak{I}$ denote the set of independent subsets of N. Let a_i denote the number of distinct independent subsets of N having i nodes. It is easy to verify that:

THEOREM 1:

The equilibrium probability that solves equation (1) is given by the following expression:

$$\pi(A) = \begin{cases} \rho^{|A|}/Z & \text{if } A \in \mathfrak{I} \\ 0 & \text{otherwise} \end{cases} \qquad (2)$$

$$\text{Where } \rho \triangleq \lambda/\mu \quad \text{and} \quad Z \triangleq \sum_{A \in \mathfrak{I}} \rho^{|A|} = \sum_i a_i \rho^i$$

The function $Z(\rho)$ is called the partition function of the interference system. The function Z provides a complete description of the distribution of transmissions among different independent sets of the interference graph. This distribution, in turn, may be used to derive the steady-state behavior of the system.

At this point, it is possible to note a striking analogy between interference systems and statistical mechanics models [RUEL 69]. Loosely speaking, the equilibrium behavior of a large scale mechanical system is described by its partition function $Z = \sum_i \exp(-\beta\varepsilon_i)$, where the summation is over all microstates of the system (enumerated by i); ε_i is the energy of the i-th microstate and $\beta = 1/kT$ where T is the absolute temperature and k is the Boltzmann constant. The equilibrium probability that the system is in the state i, is given by the Gibbs distribution: $\exp(-\beta\varepsilon_i)/Z$. All thermodynamical functions describing the system (e.g., energy, pressure, entropy) are obtained in terms of simple expressions involving the partition function and its derivatives.

The analogy to statistical mechanics may now be easily drawn: An interference system microstate is described by an independent set of nodes in the interference graph. The cardinality of an independent set corresponds to the energy of a microstate. To complete the analogy we define the temperature of an interference system as:

$$T \triangleq \frac{-1}{k \ln \rho}$$

(where k is the Boltzmann constant). Note that $\rho=0$ corresponds to $T=0$ and $\rho=1$ corresponds to $T=\infty$, so traffic increase is associated with raising the temperature. Thus temperature measures the activity level of an interference system.

With these definitions the partition function may be rewritten as:

$$Z = \sum_{A \in \mathcal{J}} e^{-|A|/kT}$$

in complete analogy to the statistical mechanics model.

Let us pursue the analogy further and interpret the "energy" associated with a microstate. Define the following pair-interaction potential function among nodes x, y of the interference graph G [RUEL 69, PRES 74]:

$$U(x,y) \triangleq \begin{cases} -\infty & \text{if } x \text{ and } y \text{ are neighbors in } G \\ (1/2)\ln\rho & \text{if } x=y \\ 0 & \text{otherwise} \end{cases}$$

This potential represents a fixed "charge" of $(1/2)\ln\rho$ associated with each transmission and an infinite repulsion among interfering transmissions. The partition function of the interference system may now be rewritten as:

$$Z = \sum_{A \subseteq N} e^{V(A)}$$

where $$V(A) \triangleq \sum_{x,y \in A} U(x,y)$$

is the respective potential of the set of nodes A. This potential is similar to that associated with a lattice interaction model of a "hard balls" gas [RUEL 69]. That is, transmissions might be thought of as balls occupying some lattice space from which other transmissions are excluded. Note also that the fixed-charge associated with a transmission may be described as $\ln\lambda - \ln\mu$, i.e., it reflects the energy gained by arrivals and lost due to departures (see [FERD 70] for an elaboration of this point). From a more abstract view [PRES 74, SPIT 71], the steady-state distribution (2) is a Markov field derived from a nearest-neighbor potential function V(A).

The significance of the representation is that at this point we have been liberated from the dynamic Markovian approach with which we started. That is, an interference system is completely described by the respective interference potential, in the sense that the steady-state distribution (and therefore any macroscopic property) are given in terms of this potential. We could have used this potential as our point of departure and avoid the birth-date Markovian model altogether. This amounts to adopting the paradigm of physicists in analyzing interference systems, rather than the paradigm of Queueing theory. To illustrate this point let us consider classical queueing systems in terms of interaction potential.

Consider first an M/M/1/m queueing system (single-server, with buffer of size m). A micro-state of the system is given by a number $0 \leq k \leq m$ indicating the length of the queue. The potential function of a microstate with k buffered customers should have a contribution of $\ln\rho$ from the "charge" of each customer and no interference terms. Therefore, the potential function is given by $V(k) = k \ln\rho$, and the respective partition function is:

$$Z = 1 + \rho + \rho^2 + \rho^3 + \ldots\ldots + \rho^m$$

The Gibbs state (distribution) associated with this partition function is easily seen to be the distribution of queue length for an M/M/1/m as given by queueing theory [KLEI 76].

Similarly, consider a Jackson network of queues [KLEI 76]. Under the assumption that no interference occurs between any two queues in the system, the partition function of the network is the product of the partition functions of the respective queues. Jackson's theorem may therefore be viewed as a statement about the existence of no interference elements in the interaction potentials between M/M/m queues connected in a network. As a matter of fact, much of the work in analysis of queueing networks can be viewed as developing methods for computing the network partition function. This also explains the inherent inability of classical Queueing theory to handle finite-buffer queueing networks: As soon as the queues start interfering with each other (e.g., by blocking) the non-interference model is no longer valid. The statistical mechanics approach admits such interference by including adequate elements in the respective potential function. While a full derivation is beyond the scope of this paper, it is possible to handle finite-buffer queueing networks by expanding the partition function in terms of the interference terms. This is similar to the virial-series expansion of interference in statistical mechanics [RUEL 69] (the first term in the series will be the Jacksonian non-interference element while higher order terms reflect mutual blocking configurations).

3. Energy, Entropy, Volume And Pressure

Statistical-mechanical analysis proceeds by identifying the potential of an interacting system, computing the respective partition function and from it deriving macroscopic thermodynamical properties of the system. Pursuing the analogy further, we are now in a position to establish a "thermodynamical" description of interference systems. We use standard definitions of thermodynamical functions (e.g., energy, free energy, entropy, pressure) in terms of the partition function [RUEL 69], and then interpret their respective significance in terms of interference systems.

The energy of a microstate has been seen to correspond to the size of the respective independent set of transmissions. The global energy thus corresponds to the average number of concurrent transmissions processed by the interference system, in other words, the thruput of the system.

More formally, the energy of an interference system is defined by the logarithmic derivative:

$$U \triangleq \rho \frac{\partial \ln Z}{\partial \rho} = (1/Z)\sum_i i a_i \rho^i$$

Entropy and pressure are defined in terms of derivatives of the free energy function. The free energy function is defined by:

$$F \triangleq \ln Z \;/\; \ln\rho$$

The entropy of an interference system is defined by:

$$S \triangleq (1/T)(U - F) = -k[(\ln\rho)U - \ln Z] \qquad (3)$$

This entropy is related to the entropy of the steady-state distribution (2) by:

$$S = -k \sum_{A\epsilon\mathcal{J}} \pi(A) \ln\pi(A)$$

It is easy to show (this is a direct result of the H-theorem for reversible Markov chains, see also [BENE 63]) that:

THEOREM 2:

Given an interference graph G and an energy U the equilibrium distribution $\pi(A)$ given by equation (2) is the unique distribution over the independent sets of G maximizing the entropy

$$S \triangleq -\sum_{A\epsilon\mathcal{J}} \pi(A) \ln\pi(A)$$

subject to the constraint

$$U = \sum_{A\epsilon\mathcal{J}} \pi(A)\; |A|.$$

The proof is a trivial exercise in constrained optimization; the term $\beta \triangleq -\ln\rho$ is the respective Lagrange multiplier.

The pressure is classically defined by:

$$P \triangleq - \frac{\partial F}{\partial V} \qquad (4)$$

Where V denotes the volume of the system. The volume is an independent external parameter (similar to temperature) that needs to be defined. How should the volume of an interference system be defined? We give an informal motivation for the definition of volume.

Pursuing the hard-balls gas analogy, consider an active node of the interference graph. Such a node may be thought of as the center of a hard ball, creating an excluded volume in which other balls may not coexist. If d denotes the average degree of a node, then the expected number of transmissions blocked by an active node is λd; we use λd to indicate the volume occupied by an active ball. The total number of such balls is $n/\lambda d$, where n denotes the total number of nodes in the graph. This number corresponds to the number of moles in a gas. If the volume of a single mole of the hard balls is kT, then the total volume is given by:

$$V \triangleq \frac{nkT}{\lambda d} = \frac{-n}{\lambda d \ln\rho}$$

With this definition of the volume, the pressure of an interference system may be computed from the definition above (4):

$$P \triangleq \frac{1}{\ln\rho} \frac{\partial \ln Z}{\partial V} = -(1/\ln\rho)\frac{\partial \ln Z}{\partial n} \frac{dn}{dV}$$

$$= \frac{\lambda d}{Z} \frac{\partial Z}{\partial n}$$

Here n, the number of nodes in the graph, is used as a continuous variable. Such approximation is, of course, only meaningful for sufficiently large interference graphs.

The above arguments are similar to those used in deriving the Van der-Waals equation of state of a gas, from a microscopic hard-balls model of interaction. As an immediate result, the following equation of state for interference systems (analogous to the Van der-Waals equation) holds:

$$\frac{PV}{n} = \frac{-1}{\ln\rho} \frac{\partial \ln Z}{\partial n} \qquad (5)$$

Finally, let us interpret the significance of pressure. What is the significance of the derivative $\frac{\partial \ln Z}{\partial n}$? Consider an interference graph whose average degree is d. Suppose G' is obtained from G by an addition of a single node, denoted x. The number of independent sets with i-nodes will grow by $a_i(G')-a_i(G)$. This number represents the number of transmission configurations having i nodes, one of which is the node x. Therefore the expression

$$\frac{\partial \ln Z}{\partial n} = (1/Z)\sum [a_i(G')-a_i(G)]\rho^i$$

represents the probability that the additional node is busy. If we multiply this derivative by λd we obtain the rate at which a transmission generated by a neighbor of the node x is blocked by this node. Therefore the pressure of an interference system represents a measure of the averaged rate of blocking experienced by transmissions. This is not asimilar to the significance of pressure in a gas.

4. Examples Of Applications

In this section we compute the partition functions and analyze the performance of a few interference networks.

4.1. A System With No Interference

Consider a system capable of producing n transmissions which do not interfere with each other. The interference graph of the system is the complement of the complete graph on n vertices. The number of independent sets of i vertices is $\alpha_i = \binom{n}{i}$. Therefore the partition function of the no-interference model is:

$$Z = (1+\rho)^n$$

The energy (thruput) of the system is given by the logarithmic derivative of Z; it is easily computed to be $U = n\rho/(1+\rho)$. Therefore the thruput per transmitting node is $\rho/(1+\rho)$ which is easily seen to be the probability that the node is busy. The entropy per node is:

$$\frac{S}{n} = -(1/1+\rho)\ \ln(1/1+\rho) - (\rho/1+\rho)\ \ln(\rho/1+\rho)$$

which is easily seen to be the entropy of the nodal idle-busy distribution. The pressure of the system is easily computed to be 0. This reflects the fact that no blocking is experienced.

4.2. A Complete Interference Model

The complete interference model is defined to be a mechanism where each transmission interferes with any other. This is a crude model of a carrier-sense bus with a negligible propagation delay. The interference graph is the complete graph K_n. Let n be the number of possible transmissions. Clearly $\alpha_1(K_n)=n$ and for $i>1$ $\alpha_i=0$. The partition function of K_n is thus given by:

$$Z = 1 + n\rho \qquad (6)$$

The energy and the pressure aren given by:

$$U = \frac{n\rho}{1+n\rho} \qquad P = \frac{\lambda(n-1)\rho}{1+n\rho}$$

The pressure is identical with the blocking probability obtained When the load on the network increases (i.e., $\rho \to \infty$), the energy (thruput) of the network approaches 1 while the pressure (blocking rate) increases to infinity.

4.3 A Tandem of Radios

A more interesting example is that of a tandem of broadcast stations. Suppose every unit on the tandem may at any time instance be involved in a transmission with one of its neighbors but not both. It is easy to derive the following recursive formulas for α_i^n, the number of transmission configurations involving i units of a tandem of n units:

$$\alpha_i^{n+1} = \alpha_i^n \quad - \alpha_{i-1}^{n-1}$$

Let Z^n denote the partition function for a tandem of length n, then the above relation implies:

$$Z^{n+1} = Z^n + \rho Z^{n-1}, \qquad Z^0=Z^1=1$$

Which may be solved to yield:

$$Z^n = \frac{1}{\Delta\xi_+^{n+1}} \quad [1 - (\xi_-/\xi_+)^{n+1}] \qquad (7)$$

Here $\xi_\pm \triangleq (-1\mp\Delta)/2\rho$ and $\Delta \triangleq \sqrt{1-4\rho}$.

Since $|\xi_+/\xi_-| = (\Delta-1)/(\Delta+1) < 1$, the second term in the expression for Z^n might be ignored for sufficiently large n. In other words, when n is sufficiently large:

$$Z^n \approx \frac{1}{\Delta\xi_+^{n+1}}$$

The energy per site may be computed to be:

$$U/n \approx \frac{1+2\rho-\Delta}{1+4\rho-\Delta}$$

As ρ approaches 1 the energy per site converges to $(1/2)-(1/2\sqrt{5})$ (this may be rewritten as $\eta/\sqrt{5}$, where η is the "golden ratio" $[\sqrt{5}-1]/2$). The term 1/2 represents the thruput per site in heavy traffic if no interference existed; the amount $1/2\sqrt{5}$ is lost due to interference.

It is possible to derive similar expressions for the pressure and entropy and then obtain the following result:

THEOREM 3:

<u>When the load ρ on the network increases to infinity, the energy per site converges to 1/2 , the entropy per site converges to 0 and the pressure converges to infinity.</u>

The interpretation of these results is that when the network is loaded with transmissions, one out of any 2 possible transmissions is active and the status of a link as a transmitting or idle link is persistent (this is reflected by the zero entropy per site). That is, the tandem is filled by a train of persistent transmissions which prevent any idle link from becoming active. Such a network may be thought of as a crystal. The tandem may crystallize in many forms, i.e., any pattern of busy transmissions where all idle links are blocked is possible, see the figure below. The significance of this phenomena is in that a long-range order may be caused through tightly coupled local interactions. Such long-range order manifest itself in packing the system with transmissions. This is similar to the hard-balls model of gas which, upon crystallization, packs the space with the balls [ZIMA 79].

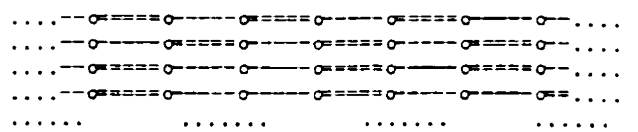

o------o An idle link
o=====o A busy link

Figure 1: Alternative Forms of Crystallization of a Tandem

4.4. A Binary Tree Switch

Consider a full binary tree switch connecting processors at the leaves. Let $n=2^{k-1}$ represent the number of leaves where k is the depth of the tree. At any time instance two processors may be involved in a transmission through their least common father on the tree. Two transmission paths which cross each other interfere with each other. Such a switch has been suggested as a multiprocessor communication mechanism [SHAW 82] and as a local area network architecture [YEMI 82]. Let us analyze the performance of the switch.

LEMMA

The number of noninterfering transmission configurations with i transmissions (on a binary tree switch with n leaves) is $a_i = \binom{n}{2i}$.

Proof:

For any transmission configuration with i transmissions corresponds a choice of 2i leaves. The converse is also true. That is, given any choice of 2i leaves, there is a unique way to associate these leaves in pairs so that the respective transmission paths do not intersect each other. Therefore the number of independent transmission sets with i transmissions is equal to the number of ways in which one can choose 2i leaves out of n.

Therefore the partition function for a binary tree switch with n leaves is given by:

$$Z = \sum_{i=0}^{\lfloor n/2 \rfloor} \binom{n}{2i} \rho^i = (1/2)[(1+\sqrt{\rho})^n + (1-\sqrt{\rho})^n]$$

From this expression the thermodynamical properties of the binary switch may be easily computed. When n is sufficiently large, Z may be approximated by:

$$Z \approx (1/2)(1+\sqrt{\rho})^n$$

Therefore, the energy per site is given by:

$$U/n \approx \frac{\sqrt{\rho}}{2(1+\sqrt{\rho})}$$

When ρ grows to infinity the energy per site converges to 1/2. That is, each leaf is involved in a transmission. It is possible to check that the entropy per site converges to 0. Again the network becomes a crystal where any two neighboring leaves are involved in a transmission while all other transmissions are inhibited.

Another interesting feature of the binary-tree switch is disclosed by examining the pressure exerted on nodes of the tree. That is, suppose we modify the definition of pressure slightly so that the average degree d is replaced by a degree of a specific node in the interference graph. The pressure at a node of the binary-tree switch is the average pressure exerted on transmissions passing through that node. It is easy to show that the pressure is maximal at the two nodes directly below the root node. This reflect the (well known) fact that the bottleneck of binary-tree switch communications is at these two nodes.

It is possible to apply similar analysis to other switching mechanisms. This will be subject of a forthcoming paper.

4.5. A Carrier Sense Packet Radio Network

Consider a spread-spectrum Packet Radio Network [KAHN 77] that uses a carrier-sense multiple access scheme. Assuming negligible propagation delays, a transmission by a node blocks its neighbors. Therefore transmissions in the network are described by an interference graph identical to the hearing graph (i.e., nodes represent radios and edges represent mutual hearing). Let λ_i denote the rate of traffic generation at node i (this includes both new arrivals as well as retransmissions) and let μ_i denote the rate of transmission duration for node i. This model of packet radio networks has been first suggested by Boorstyn and Kershenbaum [BOOR 80] and recently applied by other researchers [TOBA 82] to other models of packet radio networks (e.g., busy-tone multiple access).

Unlike the models considered so far, the traffic generation is not uniform. This generalization, however, is straightforward and requires only a redefinition of the partition function as:

$$Z = \sum_{A \in J} \prod_{i \in A} \rho_i$$

This partition function is easily seen to correspond to the following pair potential function:

$$U(x,y) \triangleq \begin{cases} -\infty & \text{if x and y are neighbors in G} \\ (1/2)\ln\rho_x & \text{if } x=y \\ 0 & \text{otherwise} \end{cases}$$

Here $\rho_i \triangleq \lambda_i/\mu_i$ is the utilization at node i.

Let r_i denote the fraction of the total traffic that is transmitted by node i, and let λ denote the total traffic generated in the network (i.e., $\lambda_i = \lambda r_i$). The local utilization parameters are linear functions of λ. Therefore, the parameter λ may play the role of a network temperature. The rest of the derivations follow suit identically to the case of uniform traffic generation. For further details and elaboration of the model and its consequences see [BOOR 80].

4.6. A Simple Model of a Database With Distributed Users

Consider a database consisting of a set of n files $D=\{1,2,\ldots,n\}$. A query of the database involves simultaneous acquisition of a subset of D. Two queries are said to interfere if the respective query sets intersect. This constitutes an interference system where queries are distributed agents attempting acquisition of a shared resource (the database). Assuming that all queries are equally likely and that queries are generated from a poisson arrival source and last an exponentially distributed time, the interference model holds. The nodes of the interference graph are subsets of D and two nodes are connected if the respective subsets intersect. Henceforth we shall use alternatively nodes, query sets and queries to indicate the same thing.

Let α_i^n denote the number of independent sets of i queries, over a database with n files.

With the above definitions it is easy to see that the following recursive relation holds:

$$\alpha_i^{n+1} = \alpha_{i-1}^n + (i+1)\alpha_i^n$$

This is a result of the following argument. Consider an independent set with i queries over n+1 files. Either the n+1-th file is a member of a query in the set or not. If the n+1-th file is not a member, then the respective set is also independent over the set of files {1,2,....,n}. If the n+1-th file is a member of a query, then by eliminating it from the query we obtain an independent set over {1,2,....n} of cardinality either i or i-1.

From the above relations we obtain the following recursive relation connecting the partition function of the database of n files with that of n+1 files.

$$Z^{n+1} = (\rho+1)Z^n + \rho\frac{\partial Z^n}{\partial\rho}$$

We define the compound transform:

$$G(s,\rho) \triangleq \sum_{n=0}^{\infty} Z^n s^n$$

to obtain the following differential equation (from the recursive relations for the partition functions).

$$0 = \frac{\partial G}{\partial\rho} + G[1 - (1-s)/\rho s] + 1/\rho s$$

From the solution of this equation one can recover the transforms Z^n and compute the thruput and pressure (blocking) of queries in the database. A detailed solution is beyond the scope of this paper and will be provided elsewhere.

5. Order-Disorder Phenomena And Critical Phase Transitions

One of the major achievements of statistical mechanics is in explaining the phenomenon of critical phase transitions. A critical phase transition occurs, for example, when a ferromagnetic metal is cooled to its Curie's temperature, at which point the metal exhibits spontaneous magnetization. Critical phenomena involve a discontinuity in some physical properties of matter when it is cooled. Such discontinuity cannot occur in a finitely small system since all physical properties depend on the temperature analytically. Therefore critical phase transitions are a direct result of the scale.

Are there critical phase transitions in interference networks? In asking this question our interest are in a broader set of questions: Are there phenomena that one may expect to exist in large-scale networks and would not be observed in small networks? Such phenomena should be of a major interest since both our intuition and analysis are usually guided by extrapolation from small scale systems to large ones. Is it possible that such extrapolation will fail? Are there phenomena involving discontinuous jump (change in phase) in the global behavior of a network as the utilization parameters increase?

In this section we show that a large class of interference models does not exhibit critical phase transitions. However, we show how a slightly different model may exhibit critical phenomena.

Consider the tandem of radios of the previous section. Interference was created among neighboring hearing links in the radio network. Therefore an independent set of transmissions corresponds to a matching in the graph of the communication network. This observation is true for a network where interference is between neighboring links.

We shall call networks where all interference is among transmissions over neighboring links, <u>matching networks</u>. For a matching network with n nodes and m edges, the matching polynomial is defined [HEIL 70, GODS 81, GUTM 82] as:

$$\alpha(x) \triangleq \sum_{k=0}^{\lfloor n/2 \rfloor} (-1)^k \alpha_k x^{n-2k}$$

(Here α_k is the number of matchings in the network with k edges.) It is easy to see that:

$$Z(\rho) = (\sqrt{\rho}/j)^n \, \alpha(j/\sqrt{\rho}) \qquad (8)$$

(Here $j=\sqrt{-1}$.) Where the partition function on the left is the one corresponding to the interference graph of the respective matching network.

A theorem due to Heilmann and Lieb [HEIL 70] (see also [GODS 81]) states that all zeros of the matching polynomial $\alpha(x)$ are real. Therefore, equation (8) implies that <u>all roots of the partition function must be real and negative</u>. Since a critical transition point (if it exists) must be a limit of zeros of the partition function as the size of the network increases, we conclude that:

THEOREM 4:

<u>There exist no critical phase transitions in interference models associated with matching networks.</u>

Are there critical transitions in interference networks? Although the above theorem demonstrates the absence of such phenomena for a large class of models, it is possible to show that critical phenomena arise in some networks. A detailed discussion of critical phenomena in computer communication networks is beyond the scope of this paper. However, the reader is referred to [YEMI 80a, YEMI 80b, YEMI 79] for example of discontinuous critical changes of phases in some packet broadcast systems.

6. Some Critical Comments

The interference model considered in this work deserves some criticism. Some assumptions could be trivially relaxed without incurring any significant change. It is possible to consider asymmetric interference, non-exclusive interference (i.e., where two transmissions are allowed to proceed concurrently with some probability), non-uniform transmission generation rates (as in the packet radio network example) etc. Other assumptions require some discussion.

The assumption that each transmission is either idle or busy does not properly account for blocked transmissions. The easy solution of this difficulty is to assume that blocked transmissions are regenerated from the same statistics as new arrivals. This is usually a good approximation used routinely in analyzing multi-access broadcast schemes. A more adequate solution will be to introduce a blocked state and account for it in the expression for the interference potential.

Another critical objection might be that the model is too static and does not account for packet motion in a network. Again, the easy answer is the independence assumption [KLEI 76], which amounts to regeneration of each packet in all nodes on its way. Given a routing mechanism the amount of traffic on each link may be determined and the analysis above can proceed using this traffic as the traffic generation rate. This avoids the need to track the motion of individual packets. A more adequate solution might be obtained by establishing a kinetic theory of packet motion in networks.

Finally, the problems of computing partition functions for more complex models are not trivial (however, we can build upon the vast set of methods used by computational physics).

7. Conclusions

We have demonstrated that statistical mechanics offers a viable and attractive alternative to Queueing theory in that it allows coarse-grained description and analysis of networks, it can handle interference, it liberates steady-state analysis from the details of the dynamics and finally, it offers physical insights to the behavior of large-scale distributed resource sharing mechanisms. This paper presents seminal work in this direction. Future work will extend the applications to more complex and more realistic models of interference, study the thermodynamics of networks in more details, import and develop computational methods to solve problems, apply the thermodynamical description of networks to design optimum routing (e.g., to minimize pressure) and flow control algorithms (e.g., control "heat" sources) and develop a kinetic theory of traffic motion in networks.

Acknowledgement

The author is grateful to Bill Hammel for his devoted reading, notes and corrections.

REFERENCES

[BARD 80] Y., Bard. Estimation of State Probabilities Using The Maximum Entropy Principle. IBM Journal of Research and Development, September, 1980.

[BENE 63] V.E., BENES. A 'Thermodynamic' Theory of Traffic in Connecting Networks. Bell System Technical Journal XLII(3):567-607, 1963.

[BOOR 80] R.R. Boorstyn and A. Kershenbaum. Throughput Analysis of Multihop Packet Radio. In Proceedings ICC. IEEE, June, 1980.

[FERD 70] A.E., Ferdinand. A Statistical Mechanical Approach to Systems Analysis. IBM Journal of Research and Development, September, 1970.

[GODS 81] C.D. Godsil and I. Gutman. On the Theory of the Matching Polynomial. Journal of Graph Theory Vol 5:137-144, 1981.

[GUTM 82] I. Gutman and F. Harary. Generalizations of the Matching Polynomial. Technical Report, Department of Mathematics University of Michigan, 1982.

[HEIL 70] O.J. Heilmann and E.H. Lieb. Monomers and Dimers. Physical Review Letters Vol 24, June 1970.

[KAHN 77] R.E. Kahn. The Organization of Computer Resources Into a Packet Radio Network. IEEE Transactions Communications Vol COM-25:169-178, 1977.

[KLEI 76] L. ,Kleinrock. Queueing Theory. J. Wiley, 1976.

[PRES 74] C.J., Preston. Gibbs States on Countable Sets. Cambridge University Press, 1974.

[RUEL 69] D., Ruelle. Statistical Mechanics: Rigorous Results. W.A., Benjamin Inc., Reading, Mass., 1969.

[SHAW 82] D.E., Shaw. The NON-VON Supercomputer. Technical Report, Computer Science Department, Columbia University, 1982.

[SHOR 78] J. E., Shore.
Derivation of Equilibrium and Time-Dependent Solutions to M/M/∞//N and M/M/∞ Queueing Systems Using Entropy Maximization.
In Proceedings NCC. AFIPS, 1978.

[SPIT 71] F. Spitzer.
Random Fields and Interacting Particle Systems.
American Math. Society, 1971.

[TOBA 82] F.A. Tobagi.
DARPA Workshop on Analytical Research of Packet Radio Networks.

[YEMI 79] Y. Yemini and L. Kleinrock.
The Capacity of Tandem Aloha and Topology Division Multiple Access.
Technical Report, Preprint, 1979.

[YEMI 80a] Y. Yemini and L. Kleinrock.
Interfering Queueing Processes in Packet Switched Broadcast Communication.
In Proceedings IFIP Congress. IFIP, Tokyo, 1980.

[YEMI 80b] Y. Yemini.
Interference Considered Useful.
In ICCC Conference Proceedings. ICCC, 1980.

[YEMI 82] Y. Yemini.
Tinkernet: Or, Is There Life Between LANs and PBXs?.
Technical Report, Computer Science Department, Columbia University, 1982.

[ZIMA 79] J.M. Ziman.
Models of Disorder.
Cambridge University Press, London, 1979.